

# **Oracle® Enterprise Manager**

Concepts

10g Release 5 (10.2.0.5)

**B31949-10**

March 2009

Oracle Enterprise Manager Concepts, 10g Release 5 (10.2.0.5)

B31949-10

Copyright © 2009, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this software or related documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of this software. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

---

---

# Contents

<b>Preface</b> .....	xi
Audience .....	xi
Documentation Accessibility .....	xi
Structure .....	xii
Related Documents .....	xiii
Updated Terminology .....	xiii
Conventions .....	xiii
<b>1 What's New in Enterprise Manager 10.2</b>	
<b>New in Enterprise Manager 10g Release 5 (10.2.0.5)</b> .....	1-1
Enterprise Manager Framework .....	1-1
Database Management .....	1-2
Middleware Management .....	1-3
Server Management .....	1-5
Applications Management .....	1-5
Lifecycle/Provisioning Management .....	1-6
Configuration Management .....	1-7
Extensibility Framework Enhancements .....	1-7
New System Management Plug-ins and Connectors .....	1-7
<b>Features of Enterprise Manager 10g</b> .....	1-8
Grid Management Features .....	1-8
Service Management Features .....	1-9
Enterprise Reporting Features .....	1-10
Database Management Features .....	1-10
Middleware Management Features .....	1-11
Collaboration Suite Management Features .....	1-12
Broader System Coverage .....	1-13
New Target Monitoring Plug-ins and Connectors .....	1-13
Oracle Business Intelligence Management Pack .....	1-14
Enhanced Diagnostics and Management Packs .....	1-14
<b>2 Introduction</b>	
<b>Quality of Service: Bridging the Business with IT Infrastructure</b> .....	2-1
Extended Application Server Management .....	2-1
<b>Lowering Management Costs Through the Grid</b> .....	2-2

Knowing What You Have and What You Are Running .....	2-3
Certified Configurations: Bringing Harmony to an Ever-Changing Data Center.....	2-3
Error-Free Automated Provisioning: Linux, Oracle RAC, Certified Images, Patches, Upgrades..	2-3
Patch Management.....	2-4
<b>Management for the Complete Oracle Grid Environment .....</b>	<b>2-5</b>
Extensions and Management Plug-ins.....	2-6
Oracle Integrated Management Partnerships .....	2-6
<b>Conclusion .....</b>	<b>2-7</b>

### 3 Overview of Enterprise Manager

<b>System Monitoring .....</b>	<b>3-1</b>
<b>Managing Groups .....</b>	<b>3-3</b>
<b>Job System .....</b>	<b>3-3</b>
<b>Information Publisher.....</b>	<b>3-4</b>
<b>Compliance Management.....</b>	<b>3-4</b>
<b>Extending Enterprise Manager .....</b>	<b>3-5</b>
<b>Managing Targets .....</b>	<b>3-6</b>

### 4 Enterprise Configuration Management

<b>Hardware and Software Configurations.....</b>	<b>4-1</b>
Collected Configurations .....	4-1
Viewing Configurations.....	4-2
Comparing Configurations.....	4-4
The Enterprise Configuration.....	4-4
Viewing the Enterprise Configuration.....	4-4
Searching the Enterprise Configuration .....	4-5
Viewing the Enterprise Configuration History .....	4-6
Client Configurations .....	4-6
Client System Analyzer in Grid Control .....	4-6
Client System Analyzer Deployed Independently .....	4-7
Accessing Configuration Pages in Grid Control .....	4-7

### 5 Database Management

<b>Introduction to Database Management .....</b>	<b>5-1</b>
Database Control Versus Grid Control.....	5-1
Database Home Page as the Management Hub .....	5-2
<b>Monitoring Databases .....</b>	<b>5-3</b>
Assessing Database Performance .....	5-3
Diagnosing Problems Interactively .....	5-5
Investigating a Wait Class .....	5-5
Viewing SQL Details .....	5-6
Viewing Session Details.....	5-7
Viewing Top Activity .....	5-8
Diagnosing Problems Automatically .....	5-8
Using Additional Diagnostic Pages.....	5-9

Top Consumers .....	5-9
Instance Activity .....	5-10
Historical SQL .....	5-10
Using Other Performance Tools.....	5-11
Topology Viewer.....	5-11
Metric Baselines.....	5-12
Memory Access Mode.....	5-13
Hang Analysis .....	5-13
<b>Administering Databases .....</b>	<b>5-14</b>
Managing Storage Objects .....	5-14
Using Database Configuration Features.....	5-14
Using Automatic Storage Management.....	5-14
Converting Single Instances to Oracle Real Application Clusters.....	5-15
Converting to a Locally-Managed Tablespace.....	5-15
Controlling Resources with Resource Manager .....	5-15
Tracking Statistics to Improve Database Performance .....	5-16
Using Oracle Scheduler .....	5-16
Working With Database Schemas.....	5-16
Managing Database Objects .....	5-16
XML Database Features in Oracle Enterprise Manager .....	5-17
Managing Users and Privileges .....	5-17
Managing Materialized Views .....	5-17
About Change Management .....	5-18
Using Advisors in Oracle Enterprise Manager.....	5-18
<b>Masking Sensitive Data for Non-production Use .....</b>	<b>5-19</b>
Security and Regulatory Compliance .....	5-19
Roles of Data Masking <i>Users</i> .....	5-20
Format Libraries and Masking Definitions .....	5-20
Recommended Data Masking Workflow .....	5-21
Data Masking Task Sequence.....	5-21
Defining New Masking Formats.....	5-23
Using Oracle-supplied Predefined Masking Formats .....	5-25
Patterns of Format Definitions.....	5-25
Category Definitions.....	5-25
Installing the DM_FMTLIB Package .....	5-26
Creating a Masking Definition.....	5-27
Adding Dependent Columns .....	5-30
Cloning the Production Database.....	5-30
Deterministic Masking .....	5-30
Use Cases.....	5-31
Related Oracle Security Offerings .....	5-32
<b>Maintaining Databases .....</b>	<b>5-32</b>
Using Backup.....	5-32
Recovery .....	5-33
Flashback Recovery.....	5-33
Backup and Recovery Settings .....	5-34
Managing Restore Points .....	5-34

Overview of Data Guard.....	5-34
<b>Additional Maintenance Features.....</b>	<b>5-35</b>
Exporting and Importing Features .....	5-35
Cloning Databases or Transporting Tablespaces .....	5-35
Overview of Oracle Streams .....	5-35
Working with Software Configurations .....	5-36
Using Database Software Patching.....	5-36
<b>Monitoring Oracle Real Application Clusters.....</b>	<b>5-37</b>
Cluster Cache Coherency .....	5-37
Cluster Interconnects .....	5-38
Cluster-Managed Database Services .....	5-39
Oracle Clusterware and High Availability.....	5-40
Database High Availability .....	5-41

## 6 Service Management

<b>Introduction to Services .....</b>	<b>6-1</b>
Defining Services in Enterprise Manager .....	6-1
<b>Modeling Services.....</b>	<b>6-2</b>
Availability.....	6-2
Performance and Usage .....	6-3
Business Metrics .....	6-3
Setting Service-Level Rules.....	6-4
Monitoring Templates for Services.....	6-4
<b>Managing Systems .....</b>	<b>6-4</b>
Creating Systems.....	6-4
Monitoring Systems .....	6-5
<b>Monitoring Services.....</b>	<b>6-6</b>
Services Dashboard.....	6-6
System Topology .....	6-7
Service Topology .....	6-8
Reports.....	6-9
Notifications, Alerts, and Baselines.....	6-9
Service Performance .....	6-9
Monitoring Web Application Services.....	6-10
Transactions.....	6-10
End-User Performance Monitoring.....	6-10
Monitoring Forms Applications.....	6-11
Forms Transactions.....	6-11
End-User Performance Monitoring.....	6-11
<b>Diagnosing Service Problems.....</b>	<b>6-12</b>
Root Cause Analysis .....	6-12
Diagnosing Web Application Problems .....	6-12
Interactive Transaction Tracing .....	6-12
Request Performance Diagnostics .....	6-13

## 7 Change Management for Databases

<b>Overview of Change Management for Databases .....</b>	<b>7-1</b>
--	------------

<b>Dictionary Baselines</b> .....	7-2
Scope Specification.....	7-2
Capturing a Dictionary Baseline Version .....	7-4
Working With A Dictionary Baseline Version.....	7-4
Working With Multiple Dictionary Baseline Versions .....	7-5
<b>Dictionary Comparisons</b> .....	7-7
Defining Dictionary Comparisons.....	7-7
Working with Dictionary Comparison Versions.....	7-9
<b>Dictionary Synchronizations</b> .....	7-11
Defining Dictionary Synchronizations.....	7-12
Creating a Synchronization Definition from a Comparison.....	7-14
Working with Dictionary Synchronization Versions.....	7-15
The Dictionary Synchronization Cycle .....	7-15
Creating Additional Synchronization Versions.....	7-21

## 8 Middleware Management

<b>Introduction to Middleware Management</b> .....	8-1
Out-of-Box Management Using Oracle Application Server Control.....	8-1
Centralized Management Using Enterprise Manager .....	8-2
<b>Monitoring Middleware Targets</b> .....	8-3
What Gets Monitored? .....	8-3
Viewing Critical Information .....	8-8
General Information .....	8-8
Alerts and Host Alerts.....	8-9
Policy Violations .....	8-9
Viewing Information About Applications Deployed .....	8-9
Viewing Members .....	8-10
Monitoring Performance.....	8-11
Using Performance Metrics .....	8-11
Automated Monitoring and Alerts .....	8-11
Viewing Topology.....	8-12
Monitoring Dashboard.....	8-14
<b>Administering Middleware Targets</b> .....	8-14
Managing and Creating Blackouts .....	8-15
Managing and Creating Jobs .....	8-15
Creating Services from J2EE or Web Applications .....	8-15
Creating Infrastructure and Aggregate Services .....	8-16
Adding SOAP Tests for Partner Links .....	8-17
Adding SOAP Tests for Oracle Service Bus Proxy Services .....	8-17
Adding SOAP Tests for Web Services of Applications Deployed to Oracle WebLogic Managed Server .....	8-18
Provisioning BPEL Processes .....	8-18
Provisioning Oracle Service Bus Resources .....	8-18
Creating Log Rotation Jobs.....	8-19
Managing Configurations .....	8-20
Extensible Monitoring .....	8-21
Diagnosing Performance Issues with Top Reports.....	8-21

Analyzing Historical Performance .....	8-22
<b>Maintaining Middleware Targets .....</b>	<b>8-22</b>
Managing Backup and Recovery of the Application Server Environment .....	8-22
Cloning and Patching the Middleware Environment .....	8-23
<b>Managing Identity Management Targets .....</b>	<b>8-23</b>
About Access and Identity Management .....	8-24
Access Manager.....	8-24
Access Server .....	8-24
Identity Server .....	8-24
Identity Manager.....	8-25
Identity Federation .....	8-25
Using Grid Control for Monitoring Identity Management Targets .....	8-25
Identity Component Server Home Page.....	8-26
Identity Management Systems.....	8-27
Configuring Identity Management Systems.....	8-27
Creating an Access Service .....	8-28
Creating an Identity Service .....	8-28
Creating an Identity Manager Service .....	8-29
Creating an Identity Federation Service .....	8-29
Monitoring Services.....	8-29
Identity Management Root Cause Analysis.....	8-30
Automated Identity Management and Alerts.....	8-30
Diagnosing Identity Management Performance and Availability Problems.....	8-30
Leveraging the Grid Control Management Framework .....	8-31

## 9 Host and Third-Party Target Management

<b>Host Management .....</b>	<b>9-1</b>
Monitoring Hosts .....	9-2
Host Home Page .....	9-2
Host Performance Page.....	9-3
Host Configuration Page .....	9-4
Operating System Details .....	9-4
Hardware Details .....	9-5
Log File Monitoring .....	9-6
Program Resource Utilization Monitoring.....	9-6
File and Directory Monitoring .....	9-7
Hardware Monitoring for Dell PowerEdge Systems.....	9-7
Storage Resource Tracking .....	9-8
Host Administration.....	9-9
Linux Host Administration .....	9-9
Execute Host Command .....	9-10
Job System .....	9-10
Remote File Editor .....	9-10
Managing Host Deployments .....	9-11
<b>Third-Party Target Management .....</b>	<b>9-11</b>
Support for Third-Party Application Servers .....	9-11
Support for Third-Party Plug-Ins .....	9-12



Network Appliance Filer .....	9-13
F5 Load Balancer Switches.....	9-14
<b>10 Oracle Collaboration Suite Management</b>	
<b>Visualizing Your Collaboration Suite Deployment</b> .....	10-1
Collaboration Suite Services .....	10-1
Service Monitoring.....	10-3
Collaboration Suite Systems .....	10-4
<b>Automated Collaboration Suite Monitoring and Alerts</b> .....	10-5
<b>Oracle Collaboration Suite Root Cause Analysis</b> .....	10-5
<b>Diagnosing Oracle Collaboration Suite Performance and Availability Problems</b> .....	10-5
<b>Collecting and Using Collaboration Suite Client Configuration Data</b> .....	10-5
<b>Oracle Collaboration Suite Reports</b> .....	10-6
<b>Leveraging the Grid Control Management Framework</b> .....	10-7
<b>Administration for Oracle Collaboration Suite Instances</b> .....	10-7
<b>11 Identity Management</b>	
<b>About Access and Identity Management</b> .....	11-1
Access Manager .....	11-2
Access Server .....	11-2
Identity Server .....	11-2
Identity Manager.....	11-2
Identity Federation.....	11-3
<b>Using Grid Control for Monitoring Identity Management Targets</b> .....	11-3
Identity Component Server Home Page.....	11-3
<b>Identity Management Systems</b> .....	11-4
Configuring Identity Management Services .....	11-5
Creating an Access Service .....	11-5
Creating an Identity Service .....	11-6
Creating an Identity Manager Service .....	11-7
Creating an Identity Federation Service .....	11-7
Monitoring Services .....	11-7
<b>Identity Management Root Cause Analysis</b> .....	11-8
<b>Automated Identity Management Monitoring and Alerts</b> .....	11-8
<b>Diagnosing Identity Management Performance and Availability Problems</b> .....	11-8
<b>Leveraging the Grid Control Management Framework</b> .....	11-8
<b>12 Oracle Beehive Management</b>	
<b>About Oracle Beehive</b> .....	12-1
<b>Using Grid Control for Monitoring Oracle Beehive Targets</b> .....	12-2
<b>Discovering Oracle Beehive Sites</b> .....	12-4
Oracle Beehive Site Home Page .....	12-5
Rediscovering an Oracle Beehive Site .....	12-6
<b>Oracle Beehive System and Services</b> .....	12-6
<b>Automated Oracle Beehive Monitoring and Alerts</b> .....	12-8

## 13 Lifecycle Management

<b>Lifecycle Management Overview</b> .....	13-1
Advantages of Using Enterprise Manager Provisioning and Patching Features .....	13-2
Wide Coverage Across the Stack .....	13-3
<b>Software Provisioning</b> .....	13-3
Deployment Procedures .....	13-3
Cloning .....	13-4
<b>Server Provisioning</b> .....	13-4
<b>Provisioning Concepts</b> .....	13-5
Components.....	13-5
Directives.....	13-5
Images.....	13-5
Software Library.....	13-6
<b>Using the Provisioning Features</b> .....	13-6
Setting up and configuring the environment.....	13-6
Creating Components, Directives, and Images .....	13-6
<b>Patching</b> .....	13-7
Patching Through Deployment Procedures.....	13-8
Patching Features .....	13-8
Application of Critical Patch Updates .....	13-10
Applying Critical Patches in Offline Mode.....	13-10
Accessing Critical Patch Advisories Pages in Grid Control .....	13-11
Application Server Patching.....	13-11
Linux Host Patching .....	13-11
Enhanced Linux Patching for ULN.....	13-12
Patching Features .....	13-12
Operating Systems Patching.....	13-13
Accessing Patching Pages in Grid Control.....	13-13

## 14 Virtualization Management

<b>Introduction to Virtualization Management</b> .....	14-1
Virtualization Central Page as the Management Hub.....	14-3
Oracle VM Server Home Page.....	14-6
Guest VM Home Page .....	14-8
<b>Monitoring Oracle VM</b> .....	14-9
<b>Administering Oracle VM</b> .....	14-9
Editing Guest VM.....	14-10
<b>Provisioning Oracle VM Server and Guest VM</b> .....	14-11
Creating Virtual Server Pool.....	14-11
Registering Virtual Servers .....	14-12
Creating Guest Virtual Machines .....	14-13
<b>Oracle Beehive Metrics</b> .....	A-1

## Glossary

## Index

---

---

# Preface

This manual introduces Oracle Enterprise Manager and describes the concepts required to use the product effectively to manage your Oracle environment.

Note that later versions of this and other Enterprise Manager books may be available on the Oracle Technology Network:

<http://www.oracle.com/technology/documentation/oem.html>

## Audience

This manual is intended for all users of Oracle Enterprise Manager.

## Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at <http://www.oracle.com/accessibility/>.

### Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

### Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

### TTY Access to Oracle Support Services

To reach AT&T Customer Assistants, dial 711 or 1.800.855.2880. An AT&T Customer Assistant will relay information between the customer and Oracle Support Services at 1.800.223.1711. Complete instructions for using the AT&T relay services are available at <http://www.consumer.att.com/relay/tty/standard2.html>. After the AT&T Customer Assistant contacts Oracle Support Services, an Oracle Support

Services engineer will handle technical issues and provide customer support according to the Oracle service request process.

## Structure

The following preface and chapters discuss in greater detail Enterprise Manager functionality and applications:

### **Chapter 1, "What's New in Enterprise Manager 10.2"**

Presents an overview of the new features in Oracle Enterprise Manager 10g Release 5 (10.2.0.5).

**Chapter 2, "Introduction"** Introduces, at a high-level, and puts into business-use context the features and benefits of using Enterprise Manager.

**Chapter 3, "Overview of Enterprise Manager"** Describes the subsystems in Enterprise Manager.

**Chapter 4, "Enterprise Configuration Management"** Explains how Enterprise Manager simplifies the monitoring and management of the deployments in your enterprise.

**Chapter 5, "Database Management"** Specifies the features Enterprise Manager includes to monitor, manage, administer, and maintain Oracle Databases.

**Chapter 6, "Service Management"** Explains how Enterprise Manager monitors and manages services and Web applications to determine the availability, performance, usage, and overall effectiveness of business implementations.

**Chapter 7, "Change Management for Databases"** Describes the Change Management feature in Enterprise Manager.

**Chapter 8, "Middleware Management"** Describes how Enterprise Manager provides an easy way to centrally manage your middle-tier Oracle Application Server environment.

**Chapter 9, "Host and Third-Party Target Management"** Describes Enterprise Manager's built-in features for managing host targets, as well as out-of-box third-party targets from non-Oracle vendors.

**Chapter 10, "Oracle Collaboration Suite Management"** Describes the new capabilities for Collaboration Suite management, including service and system monitoring, service topology, monitoring dashboards, and client data collection.

**Chapter 11, "Identity Management"** Describes how to manage your Identity Management targets.

**Chapter 12, "Oracle Beehive Management"** Describes how you can use Enterprise Manager to manage your Oracle Beehive targets.

**Chapter 13, "Lifecycle Management"** Provides an overview of the server and software provisioning and patching features introduced in Enterprise Manager 10gR2.

**Chapter 14, "Virtualization Management"** Introduces the concept of Virtualization Management in Enterprise Manager.

**Appendix A, "Oracle Beehive Service Dashboard Metrics"** Lists the Beehive metrics.

## Glossary

Provides definitions of terms specific to Enterprise Manager.

## Related Documents

For more information, see the following manuals in the Oracle Enterprise Manager 10g Release 2 documentation set:

- *Oracle Enterprise Manager Grid Control Quick Installation Guide*
- *Oracle Enterprise Manager Grid Control Installation and Basic Configuration*
- *Oracle Enterprise Manager Advanced Configuration*
- *Oracle Enterprise Manager Configuration for Oracle Collaboration Suite*
- *Oracle Enterprise Manager Policy Reference Manual*
- *Oracle Enterprise Manager Metric Reference Manual*
- *Extending Oracle Enterprise Manager*
- *Oracle Enterprise Manager Command Line Interface*
- *SNMP Support Reference Guide*

The latest versions of this and other Enterprise Manager books can be found at:

<http://www.oracle.com/technology/documentation/oem.html>

Oracle Enterprise Manager also provides extensive online help. Click **Help** on any Oracle Enterprise Manager page to display the online help system.

## Updated Terminology

In previous versions of Oracle Enterprise Manager, Service Level Management was known as either Application Service Level Management or Application Performance Management.

## Conventions

The following text conventions are used in this document:

<b>Convention</b>	<b>Meaning</b>
<b>boldface</b>	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.



---

# What's New in Enterprise Manager 10.2

Enterprise Manager 10g includes many new features that add to your ability to monitor and manage your enterprise.

## New in Enterprise Manager 10g Release 5 (10.2.0.5)

The following features and/or enhancements are available with the latest Enterprise Manager 10g Release 5 (10.2.0.5).

### Enterprise Manager Framework

- *Enhanced Monitoring Templates:* Custom monitoring settings can be specified for targets as they are added to Enterprise Manager. In addition, a new out-of-box report and reporting elements allow you to more easily review template application history operations.
- *Alert Management Enhancements:* On-demand evaluation of alerts allow administrators to quickly verify whether the fixes implemented for alerts result in clearing of the alert. In addition, administrators can better manage their log-based alerts (e.g., alert log alerts) by setting duration-based notification rules that clear such alerts on a periodic basis, or by using new EMCLI verbs that support bulk clearing of such alerts.
- *Notification Enhancements:* Repeat notifications are now supported for all notification methods – OS Command, PL/SQL procedures and SNMP traps. Administrators can customize the format of the e-mail notifications sent by Enterprise Manager to better satisfy their operational requirements. Customizations include the changing of the content and inclusion of target properties and other information that provide more business and operational context about the alert.
- *Enhanced Groups Functionality:* To streamline the management of privileges for groups and their members, you can define a new type of group called Privilege Propagating Groups. A privilege granted to a user on this type of group is automatically applied to its member targets. Group creation is also enhanced by allowing you to select member targets based on business criteria.
- *Enhancements to Enterprise Manager Backup/Recovery:* The EMCTL utility has been enhanced to support operations for the backing up and recovery of Enterprise Manager components. These include the ability to re-synchronize the repository, export and import OMS configurations, and re-synchronize an agent based on information in the repository.

- *Enhancements to Management Pack Licensing:* You can now perform bulk updates of Management Pack licenses across multiple targets simultaneously using the Enterprise Manager console or EMCLI utility.
- *Enhanced Security Features:* Enterprise Manager now provides support for the following:
  - New fine-grained target privileges have been added to support the principle of least privilege. The newly added privileges are: Blackout Target, Manage Target Metrics, Configure Target and Manage Target Alerts. These new privileges can be granted to roles and users independent of each other allowing the administrator to grant only the least privileges required.
  - Support for configuring third-party certificates for secure communication between Oracle Management Service, Oracle Management Agent, and the Web browser is now available.
  - An enhanced user interface that allows you to manage Privilege Delegation settings and support Privilege Delegation for User Defined Metrics, Corrective Actions and Database Replay features is provided.
  - Enhanced auditing capabilities including enriched audit records, audit data search, built-in externalization services to externalize audit data into external store, and compatibility with Oracle Diagnostic Logging (ODL) format to allow integration with Oracle Audit Vault are provided.

## Database Management

- *Management Support for Oracle Database 11g Release 1:* Enhanced management and monitoring support for new Oracle Database 11g features such as ADDM for RAC, Real-time SQL Monitoring, Partition Advisor, and Automatic SQL Tuning are provided.
- *Database Replay:* Comprehensive automation provides a simple, user-friendly workflow for capturing production workload, copying it to a test system, setting up software and test databases to the point time of capture, deploying replay clients, orchestrating the replay process, and analyzing the replay results.
- *Database Cloning:* Creation of database clones from existing database backups is supported.
- *Enhanced Data Masking:* Data Masking now supports the following enhancements:
  - Condition-Based Masking provides the ability to mask sensitive data in a given column with different mask formats, based on specified conditions.
  - Compound Masking supports masking a set of related columns in a table as a group, enabling sensitive data that spans multiple columns of a table to be consistently masked.
  - A secure clone and mask workflow provides the ability to apply one or more mask definitions when cloning a database.
- *Database Change Propagation:* The Dictionary Synchronization feature in the Change Management Pack supports propagation of schema changes from a Dictionary Baseline or a database to a target database.
- *Disable Alerting for UNDO, TEMP Tablespaces:* A streamlined method for disabling alerts for all UNDO and TEMP tablespaces is provided. Alerting for specific, named UNDO and TEMP tablespaces continues to be supported.



- *Enhanced Real Application Clusters Management:* Enterprise Manager provides the following Real Application Clusters (RAC) enhancements:
  - Customizable tile-based views to monitor waits and other metrics across multiple RAC nodes in a cluster are supported.
  - Improved performance workflows for cluster cache coherency, historic views, and drilldowns are available.
  - Service-centric Monitoring, which facilitates the monitoring of workflows and drilldowns for RAC services, is provided.
- *High Availability and Configuration Features:* The following new High Availability (HA) features are available:
  - A new HA Console to monitor overall HA configuration status and initiate operations is provided.
  - The Maximum Availability Architecture Configuration Advisor page allows you to evaluate the configuration and identify solutions for protection from computer, site, storage, human and data corruption failures, enabling workflows to implement Oracle Recommended solutions.
  - Automatic configuration of Oracle-recommended Maximum Availability Architecture (MAA) for databases is supported, enabling migration of databases to ASM and conversion of single instance databases to RAC with minimum downtime.
  - A Streams dashboard, along with improved monitoring of Streams configurations, allows monitoring of Streams components as well as paths for Latency and Throughput metrics.
  - Management of Oracle Secure Backup administrative server as a target and Oracle Secure Backup File System Backup/Restore and reports is supported.
- *Database Security Features:* Enterprise Manager now provides the following database security features:
  - The ability to monitor and administer Oracle Database Vault-enabled database targets is provided. This also enables the propagation of Oracle Database Vault policies such as realms and command rules from a source database to multiple Database Vault-enabled destination databases.
  - Support for monitoring and administration of database security options including Transparent Data Encryption (TDE), Oracle Label Security (OLS), and Virtual Private Database (VPD) is now available.
- *Patching of Shared Oracle Home Real Application Clusters and Clusterware:* Patching RAC Database and Clusterware on shared ORACLE\_HOMEs is supported from Enterprise Manager 10.2.0.5. The "Patch Oracle RAC - All Nodes" deployment procedure is enhanced to handle patching of Real Application Cluster databases in a shared ORACLE\_HOME setup. The new "Patch Oracle Clusterware-All Nodes" procedure is introduced to orchestrate patching of Clusterware in a shared ORACLE\_HOME setup.

## Middleware Management

- *Enhanced Support for BPEL Management:* Reports for BPEL activity statistics, BPEL pending instance statistics, and BPEL worst performing instance statistics can be generated. In addition to existing support for BPEL 10.1.3, support for BPEL 10.1.2, 10.1.3.3 and 10.1.3.4 is also provided.

- *Enhanced Application Diagnostics:* Enterprise Manager now includes Application Diagnostics for Java (AD4J), which provides enhanced Java application monitoring and diagnostics support. The product enables administrators to proactively monitor application performance and diagnose the root cause of performance problems, such as applications crashing and hanging in a production environment.
- *Enhanced Management for Oracle WebLogic Server:* WebLogic Server management support now includes the following capabilities:
  - Monitoring the performance of top Enterprise JavaBeans (EJBs) and JavaServer Pages (JSPs) in deployed applications.
  - Discovering and monitoring Web services deployed to WebLogic Server.
  - Monitoring server resources (e.g. data sources, JMS servers, resource adapters, JOLT connection pools).
  - Viewing, comparing and tracking configuration items such as JVM vendor/version, additional tuning parameters, cluster configuration, JMS resources, virtual hosts, JOLT connection pools and configuration files.
- *Oracle Coherence Management:* The Management Pack for Oracle Coherence provides comprehensive tools for managing and monitoring the entire Oracle Coherence Cluster. It covers key management areas such as discovery, monitoring, reporting, events management, configuration management, lifecycle management and process automation.
  - Automated deployment of new Coherence Clusters or addition of nodes to existing Coherence clusters is supported.
  - Monitoring of the entire Coherence cluster as a single target and identifies potential hotspots is supported.
  - Modeling and management of the entire Coherence infrastructure, including WebLogic Server, Coherence Cluster and database as a single system, is provided.
  - Performance monitoring of Coherence nodes, caches and services as well as Connection Managers and connections is provided.
  - Real-time propagation of configuration changes to nodes, caches or an entire cluster is supported.
- *Oracle Service Bus Support:* New features for deploying and managing Oracle Service Bus (OSB) instances are available:
  - Proxy and business service monitoring, along with performance monitoring functionality for OSB 10gR3 stage and action level metrics, is provided.
  - Collection of configuration information Oracle Service Bus targets is now supported. The configuration information of these targets can be refreshed, saved or compared with another target. Different versions of the same OSB can also be compared.
  - Automatic system and service creation for OSB infrastructure and OSB proxy services is supported.
  - A new procedure for deploying OSB projects and resources from the software library or by exporting them from an existing OSB target is supported.
- *Adapter Metrics:* Throughput and error metrics for different adapters are provided in graphical format.

- *Composite Application Monitor and Modeler*: Enterprise Manager now includes Composite Application Monitor and Modeler, a model-driven intelligent monitoring and modeling platform for JavaEE and SOA applications.

## Server Management

- *Oracle Virtualization (Oracle VM) Management*: Virtualization management involves the monitoring, administration, and maintenance of the virtual servers and guest virtual machines in your enterprise. Enterprise Manager is now able to monitor virtualized entities alongside the physical infrastructure and perform complete lifecycle management of virtual servers and software running on them. New features include:
  - A complete set of integrated features for managing Oracle VM.
  - Increased manageability of virtualized infrastructure and reduced virtual server sprawl due to configuration collection and tracking.
  - Performance and availability monitoring for both Oracle VM virtual servers and physical servers from a single console.
  - Faster deployment of Oracle VM Guest virtual machines, enabling the creation of multiple virtual machines, then cloning and configuring them.
  - Failover support based on proven Oracle Clusterware technology.
- *Management of RPM Repositories and Channels*: The following Linux management features have been added:
  - Enterprise Manager integrates with Unbreakable Linux Network (ULN) and provides automatic patch compliance with the latest updates made available by Oracle.
  - RPM repositories can be created from within EM console and automatically refreshed.
  - Visualization and search for packages in various channels is supported.
  - Support for creating, copying and managing custom channels is available.
  - Security advisories and other errata information made available by Oracle is also displayed.
- *Enhanced Bare Metal Provisioning*: The following enhancements for Linux bare metal provisioning have been added:
  - Multiple default images can be provisioned on the same subnet.
  - Automated DHCP configuration is supported.
  - Post-provisioning steps, including user creation, are supported.

## Applications Management

- *Enhanced Siebel Support*: The new version of Application Management Pack for Siebel adds Siebel 8.1.1 certification, Siebel workflow process and policy monitoring, Siebel event log analysis, improved Siebel HI service testing and improved application discovery. Siebel workflow process and policy monitoring enables monitoring of the individual instances of Siebel workflow processes and policies, sending alerts when these fail or are slow to execute. Event log analysis facilitates problem resolution. The Siebel HI service test now includes step level timing measurement details, the execution of key business flows to be monitored with greater granularity.

- *Real User Experience Insight*: This product maximizes the value of your business critical applications by delivering insight into real end-user experiences through reporting on user traffic requested by, and generated from, your network. It integrates performance analysis and usage analysis into a single offering, enabling business and IT stakeholders to develop a shared understanding into their application users experience. Out-of-the-box support for Siebel and E-Business Suite is provided.
- *Real User Experience Insight Accelerator for Siebel*: This feature enables out-of-the-box monitoring of Siebel applications, including user session tracking, discovery of end-user performance issues, and identification of Siebel application issues.
- *Real User Experience Insight Accelerator for E-Business Suite*: This feature enables out-of-the-box monitoring of EBS applications, including user session tracking, discovery of end-user performance issues, and identification of Siebel application issues.
- *Application Testing Suite*: This Application Quality Management (AQM) solution provides automated functional testing, load testing, and test management for Web, packaged and SOA-based applications. It is an integrated, full lifecycle solution that enables you to define and manage your application testing process, validate application functionality, and ensure that applications will perform under load.
- *Oracle Beehive Management*: Oracle Beehive services can be managed from the overview level to the individual component level. Oracle Beehive unifies common yet disjointed collaborative applications, such as e-mail, time management, instant messaging etc., and delivers them through standard protocol clients, integrated Web Services, and clients such as Microsoft Outlook.

## Lifecycle/Provisioning Management

- *Software Library Management*: A number of Software Library management features are provided:
  - Purging of Software Library entities (Components/Images/Directives/Suites/Networks) from the physical directory.
  - Pre-requisite check for Software Library space, and accessibility from multiple Oracle Management Service instances.
  - Export and import of Software Library entities such as Components/Images/Directives/Suites/Networks is supported. The exchange happens in the form of Provisioning Archive (.par) files
- *New Forms and Reports Provisioning Deployment Procedure*: Cloning an already installed and discovered Forms and Reports release 9.0.4.x onto selected hosts is supported.
- *Database Provisioning Features*: The single instance database provisioning procedure has been enhanced to handle the following:
  - NFS mounted staging locations as a source for gold images. This can be used for Provisioning multiple ORACLE\_HOMEs in low bandwidth networks.
  - Additional options to customize the software deployment.
  - Improved database creation options to use user-specified templates and response files.

- *Separation of Prerequisite Checks for Scaling Up RAC Clusters:* The "One Click Extend Cluster Database" deployment procedure has been enhanced to support prerequisite checks independent of the scale-up operation. This enables users to ensure readiness of the new machine RAC is being extended to.
- *Delete Entire Node or Cluster With Deployment Procedure:* The "Delete/Scale down Oracle Real Application Clusters" procedure has been enhanced to handle the following cases:
  - Scale down or delete partially configured nodes of a cluster. This is particularly useful to rollback a failed attempt to scale up a node.
  - Delete the entire RAC cluster setup, including partial failed deployments.

## Configuration Management

- *Compliance Management:* The following compliance management features are provided:
  - A simplified User Defined Policy interface is provided, allowing system vulnerabilities to be mitigated by defining and implementing configuration policies based on operational best practices, governance and industry standard requirements. The new interface allows you to create, edit, test, delete, export and import User Defined Policies and Policy Groups.
  - Tracking, analyzing, reporting and viewing details related to IT compliance using a single interface is supported. A rich set of out-of-the-box reports and capabilities for filtering information and flexible search capabilities across the enterprise are included.
  - Configuration Change Console, which provides real-time detection of any configuration changes made to the environment, is provided.
- *Asset Discovery and Tracking:* Configuration information is automatically collected at regular intervals and stored in the repository. This information is then available in Enterprise Manager from a central console providing a comprehensive view of all heterogeneous components in the data center.

## Extensibility Framework Enhancements

The Enterprise Manager Extensibility framework allows you to build custom plug-ins that can extend Enterprise Manager to manage and monitor applications that are not managed by Enterprise Manager out-of-the-box. Some enhancements in the plug-in framework are outlined below.

- *Customization:* The ability to customize a plug-in's homepage by adding your key charts or graphs is supported.
- *Improved Plug-in Report Navigation:* URL links can be inserted into plug-in reports, enabling ease of navigation in plug-in reporting elements.

## New System Management Plug-ins and Connectors

### New Target Monitoring Plug-ins

Enterprise Manager's target monitoring capabilities have been expanded through the addition of the following plug-ins.

- System Monitoring Plug-in for Dell Change Automation
- System Monitoring Plug-in for Oracle Exadata Storage Server

- System Monitoring Plug-in for VMware ESX Server

### **New Management Connectors**

In addition to generating tickets for alerts, the helpdesk connectors now support the generation of tickets for job status events. Also, target properties can now be used in ticket templates to provide better operational and business context to generated helpdesk tickets.

- *Helpdesk Connectors*
  - Management Connector for BMC Remedy Service Desk
  - Management Connector for HP ServiceCenter
  - Management Connector for HP Service Manager
- *Event Connectors*
  - Management Connector for HP OpenView Operations
  - Management Connector for IBM Tivoli Enterprise Console

## **Features of Enterprise Manager 10g**

The following features and/or enhancements have been available with preceding releases of Enterprise Manager 10g (10.2.0.4 and lower.)

### **Grid Management Features**

- *System Monitoring*: Lights-out, comprehensive monitoring of all components of the Oracle ecosystem. Alert handling flexibility is improved through repeating notifications. Provides support for evaluating alerts on demand (on-demand metric collection). On-demand metric collection allows you to reevaluate an alert condition after an issue has been resolved, allowing Enterprise Manager to clear the alert. Automatically apply custom monitoring settings to newly discovered targets via monitoring templates.
- *Privilege Delegation Providers*: Administrative security and flexibility are enhanced through support for privilege delegation providers for executing jobs and patching targets from Enterprise Manager. A privilege delegation provider allows a logged in administrator using his account to be authorized to execute a specific action based on a subset of privileges delegated from another account. Super Administrators can now set up regular administrator accounts to leverage the two supported privilege delegation providers: Sudo and PowerBroker.
- *Enterprise Manager Audit Framework*: Enhanced auditing capability increases tracking flexibility and utility by allowing Enterprise Manager administrators to define custom audit attributes in addition to the standard "out-of-box" attributes. Users can now enter text strings such as a work order number before performing actions in Enterprise Manager in order to have that custom string associated with all the audit records associated with that action, creating a richer audit trail.
- *Compare Monitoring Templates*: See at a glance how metric and policy settings defined in a template differ from those defined on destination targets.
- *Bare-Metal Provisioning*: Automatically deploy new fully configured systems, including operating systems and software images. Clone existing installations and software images (such as Oracle Cluster Ready Services (CRS), Oracle Real Application Clusters (RAC), and Oracle Application Server).

- *Oracle Real Application Clusters Provisioning and Conversion:* Automatically configure Oracle RAC database instances to dynamically increase database capacity to meet your needs; convert non-RAC databases to Oracle RAC databases.
- *Cloning:* Simplify deployment of Oracle solutions by cloning fully patched and tested Oracle homes to multiple hosts; clone database instances to create backups of important databases; create new single-node Oracle RAC by cloning an existing RAC Oracle home.
- *Lights-Out Patching:* Quickly identify available patches for the components in your enterprise. Find out which have not been applied and which are critical. Bring those components up to the latest patch level with just a few clicks.
- *Policy-Based Monitoring:* Apply consistent, proactive management practices across your enterprise using policies and management templates.
- *Policy Groups:* Policy groups are standards by which targets are measured. The rules within a policy group are managed as a set; that is, they are viewed, maintained, evaluated, and reported in the context of a group. In contrast, policy rules are not evaluated as a set. Policy rules are viewed, maintained, and evaluated as standalone entities. Oracle provides the following policy groups: Secure Configuration for Oracle Database, Secure Configuration for Oracle Real Application Cluster, and Secure Configuration for Oracle Listener.
- *User-Defined Policies and User-Defined Policy Groups:* You can now define policies to monitor needs specific to your organization. By using the Create Policy wizard, you can easily create policies that meet the demands of security and audit compliance requirements. You can also create your own policy groups that better meet your needs for best practices, data privacy, and configuration requirements.
- *Security Assessment:* Automatically assess the security vulnerabilities in your enterprise using the Critical Patch Facility, which provides notifications and a direct link to Oracle Support. Enhanced security rules harden the infrastructure database and host. Violation score reporting helps administrators prioritize security operations.
- *Job System:* More powerful, feature-rich job system allows you to schedule and run multi-task jobs across multiple targets, as well as receive notifications.
- *Group Management:* Enhanced group management features enable administrators to manage many components as one. Group pages roll up critical management information, such as overall status and group status history, open alerts and alert history, open policy violations and policy trend overview charts, and out-of-box group performance charts. Simplifies common administrative tasks, such as running jobs, executing custom SQL or host commands, and scheduling maintenance windows. A group "dashboard" enables proactive monitoring of status and alerts across all group members as they occur.

## Service Management Features

- *Services:* Define services to represent the business functions or applications that run in your enterprise. Measure the performance and availability of these services, receive alerts when there is a problem, identify common issues, and diagnose causes of failures.
- *Service and System Topologies:* Easily define, visualize, and manage key business functions and the assets that they rely upon.

- *Service Level Management*: Allows you to define service-level goals for availability and performance and more quickly resolve problems through rules, reports, and historical monitoring.
- *Dashboards*: Offer "at-a-glance" monitoring of all critical indicators for your services and systems.
- *Root Cause Analysis*: Diagnose service problems quickly using the Root Cause Analysis engine.

## Enterprise Reporting Features

- *Out-of-Box Reports*: Easy-access out-of-box reports for all Oracle components; reports on configuration, performance, service level, and audit.
- *Graphical Report Creation*: Easy-to-use tool to create custom reports.
- *Complete Repository Access*: Create reports using any data collected and stored in the Management Repository.
- *Report Publishing*: Proactively deliver specialized management information to users, managers, and executives using out-of-box reports, custom reports, and secure publishing features, with views accessible by any reporting product.

## Database Management Features

- *Oracle Real Application Clusters Management*: Enhanced monitoring of Oracle RAC databases and clusters to include Interconnect monitoring, cluster cache coherency diagnostics, CRS monitoring, a graphical topology view; provision Oracle RAC deployments, convert single instances, and add nodes to existing clusters.
- *Transportable Tablespaces*: Allows an administrator to move tablespaces between databases.
- *Fast-Start Failover*: Allows Data Guard to rapidly and automatically fail over to the standby database without requiring manual intervention, significantly improving the degree of availability as well as the disaster resilience for the system infrastructure.
- *Restore Points*: Guaranteed recovery of the database to a particular point in time when used with Flashback Database or Flashback Table.
- *Enhanced Backup Management*: Provides a single view of all backups run against the databases in a domain.
- *Access Advisor Enhancements*: Enhanced to expose more powerful features of the Access Advisor, including template support, tuning of a single statement, journaling, recommendation implementation status, and action implementation status.
- *Automatic Segment Advisor*: Proactively identifies segments that have significant wasted space due to data fragmentation, which are good candidates for Online Segment Shrink. You no longer need to run Segment Advisor to determine what segments have data fragmentation.
- *Online Segment Shrink Enhancements*: Extends Online Segment Shrink capability to large object (LOB) segments and index-organized table (IOT) overflow segments. Online Segments Shrink capability is now extended to all types of segments, further reducing the need to perform other more advanced reorganization methods such as Online Redefinition.



- *Oracle Streams*: Share message information within a database or from one database to another. Control what information is put into a stream, how the stream flows from database to database, what happens to messages in the stream, and how the stream terminates.
- *Adaptive Alert Thresholds*: Computes statistical alert thresholds using either static (user-defined) or dynamic (self-adjusting) baselines. Dynamic statistical baselines can significantly improve the accuracy of performance alerting, while reducing exposure to the "false positives" commonly incurred under fixed threshold schemes.

## Middleware Management Features

- *Extended Middle Tier Management*: Allows administrators to discover Oracle Application Server Farms, Oracle Application Server Clusters, Oracle HTTP Server High Availability Groups and OC4J High Availability Groups as composite targets and perform management and administration operations on them. Also supports discovery, management, and administration of Oracle BPEL Process Manager targets and Oracle BPEL Processes. Allows creation of infrastructure services and aggregate services. Also allows viewing of partner links associated with a BPEL process and adding of SOAP tests for them to test their availability.
- *BPEL Process Provisioning Deployment Procedure*: Supports deployment of BPEL processes on selected BPEL Process Managers. The BPEL process suitcase files are placed as generic components in the software library, and once it is available for deployment, the deployment procedure can be run to deploy the processes on BPEL Process Managers.
- *Oracle Service Bus (OSB) Provisioning Deployment Procedure*: Supports deployment of project resources from a source Oracle Service Bus to a target Oracle Service Bus domain. Also supports movement of configuration from one environment to another. The project resources can be exported at project level or resource level, and can be deployed either directly from the source Oracle Service Bus domain or from the Software Library if the resources are stored as generic components. Also supports use of use customization file to customize the environment variables in the changed (target) environment.
- *Oracle BI Suite Enterprise Edition Management*: Supports discovery, management, and administration of Oracle BI Suite Enterprise Edition components, mainly Oracle BI Presentation Server, Oracle BI Cluster Controller, Oracle BI Analytics Server, and Oracle BI Scheduler. Also supports discovery and monitoring of Data warehouse Application Console Server (DAC Server), which is used to create and schedule ETL task flows for Informatica server.
- *Log File Rotation Jobs*: Supports scheduled rotation of log files of OracleAS components to a default or user-defined location. Allows adding of a log rotation job as part of a multi-task job and saving of this job to the library.
- *MQ Series Support*: Supports discovery and monitoring of IBM WebSphere MQ Queue Managers and IBM WebSphere MQ Clusters. Also allows monitoring of their logical objects such as Queues and Channels.
- *JBoss Application Server Management*: Supports discovery, management, and administration of JBoss Application Servers and JBoss Partitions. Also allows refreshing of JBoss Partition membership to reflect the current deployment state.
- *Microsoft Exchange Server Management*: Supports discovery, management, and administration of Microsoft Exchange Organizations, Microsoft Exchange Routing Groups, and Microsoft Exchange Servers.

- *Automatic Creation of System and Services:* Supports automatic creation of a system and service from a J2EE or Web application after checking for dependencies. Allows creation of a service with a new or existing system based on the dependency mapping available. Also allows update of an existing service and system.
- *Configuration Management:* Extended configuration management support for OracleAS deployments. Collect, search, track, compare, and report on all Oracle Application Server hardware and software installation configurations. The product also maintains a central inventory of software configuration data of core components for all Oracle Application Server instances.  
  
As part of the configuration management tasks, you can also view the BPEL Processes, its different versions, and the suitcase files associated with each version. You can also compare the BPEL Process suitcase files of different versions and track the changes that were made to a version. This allows you to track changes and identify the cause for changed behavior of the BPEL Process.
- *Topology Views:* Provides graphical topology view of AS Farms, Clusters, and ASs, with current status, alerts and policy violations, and performance metrics. View by host, by request routing overview, or request routing details.
- *Custom Aggregate Monitoring:* Define monitoring criteria for OracleAS Farms and Clusters using customizable aggregate metrics based on the statistical aggregation of individual metrics of cluster members. View aggregate metrics as an overlay of “top best,” “top worst,” or “peak threshold” values for individual metrics.
- *J2EE Application Monitoring:* Hierarchical J2EE application view at the OracleAS cluster level shows performance across the OC4J instances on which it is deployed. Ability to plug customer application JMX MBeans into Grid Control.
- *Backup and Restore:* Supports full and incremental online backup of middle tier installations. Automatically recover instances to the last full backup, then apply incremental backups.
- *OPMN Jobs:* Supports immediate or scheduled full process control (start, stop, restart) of OPMN managed components from Grid Control.
- *Provisioning “Gold” Images:* Save standard reference images of application server targets to the central Software Library and deploy to target locations. Standardize the deployments of the application server targets in your enterprise. Create and extend OracleAS Cluster environments by creating new clusters or adding nodes to existing clusters.
- *Cloning:* Extended cloning support for different types of middle-tier installations, including WebCache, J2EE and WebCache, Portal and Wireless, Forms and Reports Services, and Business Intelligence and Forms.

## Collaboration Suite Management Features

- *Centralized Management:* Single target subtab for Collaboration Suite from which you can view your Collaboration Suite deployment(s) and navigate to specific components.
- *Automatic Discovery:* Collaboration Suite server targets, such as SMTP servers and Calendar servers, are automatically discovered and monitored in Grid Control.
- *Service Monitoring:* In addition to monitoring server components, Grid Control also monitors Collaboration Suite services from the end-user perspective.

- *Custom Service Modeling*: Fast setup of services using a simple customization user interface.
- *Topology Viewer*: Service topology view of your Collaboration Suite deployment showing relationships between different services in Collaboration Suite.
- *Root Cause Analysis*: Detects the potential root causes of any service outage.
- *Dashboards*: Services Dashboard for monitoring service availability, performance, usage, and service-level compliance. System Dashboard for monitoring exceptions in your Collaboration Suite system.
- *Client System Analyzer*: Collects Windows client data and checks for minimum Collaboration Suite requirements.

## Broader System Coverage

- *Host Management*: Out-of-box operating system, hardware, and software monitoring of the hosts on which your applications and services reside. Increased coverage for hardware monitoring, process and file-level monitoring, remote command shells, and remote file editing.
- *Linux Host Administration*: Host administration features for RedHat and SuSE Linux provide you the ability to administer the most important host functions easily. Using the centralized console, it is faster and less error prone to set up a server or troubleshoot problems. Operators with the appropriate privileges can set up and administer System Services, for example, start and stop individual services, and Network Setup, for example, NFS client configuration.
- *Enterprise Storage Reports*: Reports on allocation, utilization, roll-ups by location, line of business, application, and vendor.
- *Support for Network Devices and Storage*: Monitor network and storage devices, such as F5 Server Load Balancers and NetApp Filer network storage.
- *Support for Non-Oracle Middleware*: Support for discovery, performance monitoring, reporting, alerts, process control, and synthetic transactions for third-party application servers and application server clusters, such as IBM WebSphere and Oracle WebLogic.
- *External Product Integration*: Integrate external products using command line interface, notifications, and bi-directional messages.
- *E-Business Suite Management*: Plug-in available to enable Grid Control to manage the Oracle Applications Manager, part of the E-Business Suite.

## New Target Monitoring Plug-ins and Connectors

### New Target Monitoring Plug-ins

Enterprise Manager's target monitoring capabilities have been expanded through the addition of the following plug-ins.

- System Monitoring Plug-in for Hosts
- System Monitoring Plug-in for Non-Oracle Databases
- System Monitoring Plug-in for Non-Oracle Middleware
- System Monitoring Plug-in for Network Devices
- System Monitoring Plug-in for Storage

### **New Management Connectors**

The following management connectors are now available with Enterprise Manager.

- Management Connector for BMC Remedy Service Desk
- Management Connector for HP OpenView Operations
- Management Connector for HP ServiceCenter
- Management Connector for HP Service Manager
- Management Connector for Remedy Management Platforms

## **Oracle Business Intelligence Management Pack**

Oracle Business Intelligence Management Pack is an integrated solution for ensuring the performance and availability of Oracle Business Intelligence Enterprise Edition. Using the pack, you can do the following:

- Receive proactive alerts on impending problems on the application.
- Examine run-time statistics to troubleshoot and tune performance.
- Compare configurations across different environments.
- Keep track of configuration changes.
- Proactively enforce configuration policies to avoid configuration-related problems.

You can use the pack with other Oracle Enterprise Manager management packs and plug-ins to achieve comprehensive management of the entire application environment, including both Oracle and non-Oracle databases, operating systems, storage, and network devices.

Key product features include:

- Application performance management
  - Dashboards and reports
  - Alert notifications
  - Historical metrics
  - Blackouts
  - O/S management
- Configuration management
  - Automatic discovery
  - Configuration snapshots
  - Configuration comparisons
  - Change history
  - Configuration policies

## **Enhanced Diagnostics and Management Packs**

### **Enhanced Oracle Diagnostics Pack for Non-Oracle Middleware**

Oracle Enterprise Manager 10g introduces the Diagnostics Pack for Non-Oracle Middleware. This pack provides the following features:

- Discovery of non-Oracle middleware (Web Logic, Web Sphere, and JBoss)

- System monitoring of key performance, load, and usage metrics
- Out-of-box thresholds, events, notifications, monitoring templates, and automated corrective actions
- Real-time and historical monitoring

### **Enhanced Service Level Management Pack**

The Service Level Management Pack has added the following support:

- Recording Forms transactions for applications based on Forms Server 6i.
- DHTML recorder to record Web transactions for applications with more dynamic content.

### **Enhanced Application Management Pack for Siebel**

Oracle Application Management Pack for Siebel is a complete and integrated solution for ensuring the performance and availability of Siebel applications.

- *Siebel Enterprises and Siebel Servers*: Supports discovery and monitoring of Siebel Enterprises and Siebel Servers.
- *Siebel Workflow Process and Workflow Component Groups*: Supports discovery and monitoring of Siebel Workflow Processes and Siebel Workflow Component Groups.
- *Siebel Application Services*: Supports discovery and monitoring of Siebel applications deployed to different servers.
- *Alerts*: Provides proactive alerts on impending problems on the application. You can examine run-time statistics to troubleshoot and tune performance.
- *Event Log Monitoring*: Allows you to view a list of event logs associated with user sessions.
- *Configuration Management*: Keeps track of configuration changes and allows you to compare configurations across different environments. Proactively enforces configuration policies to avoid configuration-related problems.
- *Service Level Management*: Monitor and report on service levels delivered by the application and performance diagnostics of user and server performance.
- *Transaction Recording*: Supports recording of Siebel transactions and grouping of related steps into separate, logical entities.
- *Integration with Siebel Transaction Diagnostic*: This pack adds integrated Siebel Transaction Diagnostic support using data captured through the Siebel Application Response Measurement (SARM) framework. With this tool, you can:
  - Analyze request processing for individual user requests.
  - Analyze CPU and memory consumption of processing requests for different Siebel Server Components.
  - Graphically visualize SARM metric data
  - Compare performance profiles over time
  - Create performance reports and share them with your co-workers for collaborative performance diagnostics.

The Siebel Transaction Diagnostic tool also adds support for Siebel 7.7 and 7.8 besides being compatible with Siebel 8.0. It is fully integrated with Oracle Enterprise Manager, which eliminates the need for separate installation and user

maintenance, and it runs on all operating system platforms that Oracle Enterprise Manager supports.

---

---

## Introduction

The Oracle Grid offers a proven solution that allows businesses to heighten application performance and deliver unparalleled IT infrastructure reliability. Businesses of all sizes attain measurable IT cost savings by extending the availability of computing resources, cultivating more productive and effective administrators, and achieving higher quality of service—all direct benefits of the Oracle Grid. A critical ingredient to achieving the full benefits of the Oracle Grid is employing the right solution for managing the grid systems and, most importantly, the applications that run on those systems.

Oracle Enterprise Manager 10g is Oracle's single, integrated solution for managing all aspects of the Oracle Grid and the applications running on it. Grid Control couples a potent, top-down monitoring approach to delivering the highest quality of service for applications with a cost-effective automated configuration management, provisioning, and administration solution. This powerful combination provides unequalled management for any size Oracle data center.

This chapter provides an overview of Enterprise Manager 10g Release 2 (10.2), its rich functionality, and the unique benefits it offers to Oracle customers.

### Quality of Service: Bridging the Business with IT Infrastructure

Until recently, business success measurements and IT infrastructure performance were rarely mentioned in the same conversation. Corporations treated the two distinctly—in part because of traditional organization structures—but mostly because existing management tools offered point solutions that only addressed a portion of the larger need. IT management priorities focused primarily on maintaining individual IT systems performance, largely disconnected from real business objectives and needs.

Now companies realize that in order to manage bottom line profits and differentiate the quality of service they offer, they must gain insight into how businesses are affected by IT performance. This requirement has mandated innovation in management solutions, where tools must go beyond discrete systems management to focus on delivering application availability and performance levels that align with business priorities. Few management solutions—because of their limited view of the application and its life cycle—are up to the task.

### Extended Application Server Management

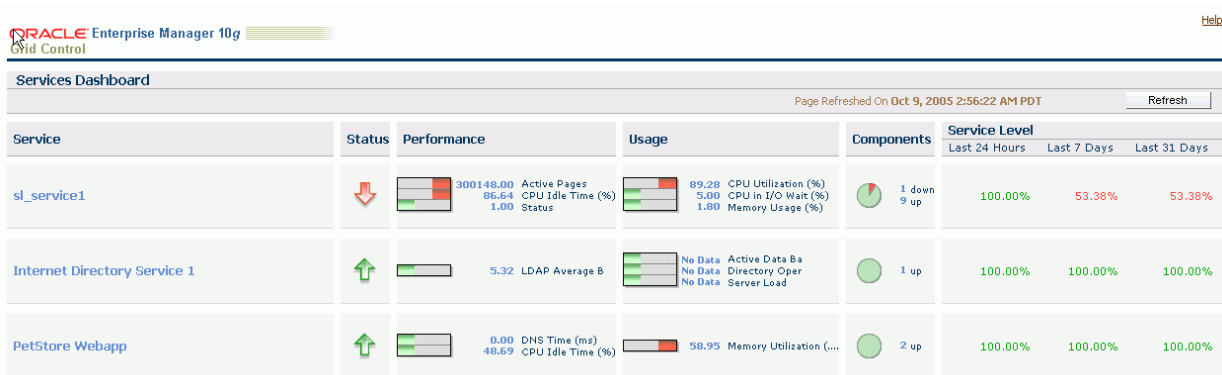
Oracle Enterprise Manager 10g, a critical element of the Oracle Grid platform, manages service level performance from a business perspective. By leveraging its "best-of-breed" system monitoring solution—including out-of-box support for databases, application servers, hosts, firewalls, load balancers, and storage—it provides a top-down approach to managing applications. This solution is unique

because of its deep system-specific coverage and the range of applications supported. No other solution offers the same level of depth, integration, and precision for Oracle customers. New functionality in Grid Control Release 2 extends the Release 1 application service level management solution with rich graphical modeling and monitoring, crisp and workflow-driven dashboards, and root-cause diagnostics, which take administrators from an end-user application availability issue through the entire application stack to pinpoint the offending IT infrastructure component. Additionally, complementary functionality allows administrators to quickly identify and triage impacted applications when low-level system issues surface.

Using the new features available in Grid Control Release 2, administrators can model and relate system components within the IT infrastructure that provide the backbone for service offerings. Leveraging these system components, they can define applications, or services, that depend upon those systems, and select business transactions that are measured to determine availability and performance of those services. Grid Control supports over 20 different transaction protocols, including custom scripts, dramatically extending the breadth of coverage to meet the application monitoring needs of nearly every environment.

Using Grid Control’s graphical monitoring dashboards, administrators, managers, and business owners gain real-time understanding of the status of the business’s most important services in a single view. Monitoring dashboards provide top-level views of individual or groups of services, as well as system-level views describing the relationships and status of IT system components. Brought together with newly introduced root-cause availability management capabilities, dashboards dramatically reduce the time it takes to resolve issues that directly impact performance. Service and system "dashboards" are available out-of-box for Oracle packaged applications.

Figure 2–1 Consolidated Dashboard View of Multiple Services



## Lowering Management Costs Through the Grid

The number of systems and diversity of services in the Grid can be small or quite large. Historically, the cost of managing large sets of systems has increased linearly, or more, with each new system added to the enterprise. It's simply not possible to maintain management costs in the Grid environment with traditional management solutions that do not scale. The fundamental value proposition of Grid Control is to maintain management costs at a flat or near flat rate as the number of managed systems increases. There are a number of things that make this possible.



## Knowing What You Have and What You Are Running

For years, IT departments have relied upon the knowledge of key individuals who manually installed, configured, and maintained the systems running throughout an enterprise. Keeping systems up and running was often a finely-tuned recipe book of details written in a language only those closely involved could understand. Yet, as distributed systems evolved and service architectures became more complex and heterogeneous, manual solutions were no longer viable from a cost-effective standpoint—and more pressingly, from a quality of service standpoint.

Administrators are now responsible for far more systems, and the relationships between those systems are much too complicated to track manually. Firewalls, switches, load-balancers, application servers, databases, and clusters are all distributed and connected through complex rules that demand reliable, automated configuration management.

Grid Control's configuration management solution is designed to eliminate the costs and errors that plague traditional manual configuration management solutions. Grid Control centrally tracks hardware, software installation (including patch levels), and software configuration data for every service and system that it manages. This detailed information is regularly and automatically collected and updated as changes occur. Central storage of enterprise configuration information lays the foundation for defining, deploying, auditing, enforcing, and maintaining the systems throughout the Grid.

## Certified Configurations: Bringing Harmony to an Ever-Changing Data Center

IT professionals will agree—the fewer variations in the types of systems deployed and the configurations used to deploy them, the easier systems are to manage and the more reliable they will be. That said, it is impossible to expect, even with the most mature business practices, that all systems can follow the same standards of deployment. Using a robust configuration management solution, administrators can rely upon automated tools to ensure that all systems deployed follow specified practices and rules. This way, only pre-tested, pre-certified configurations can enter the production "food chain" of a busy data center.

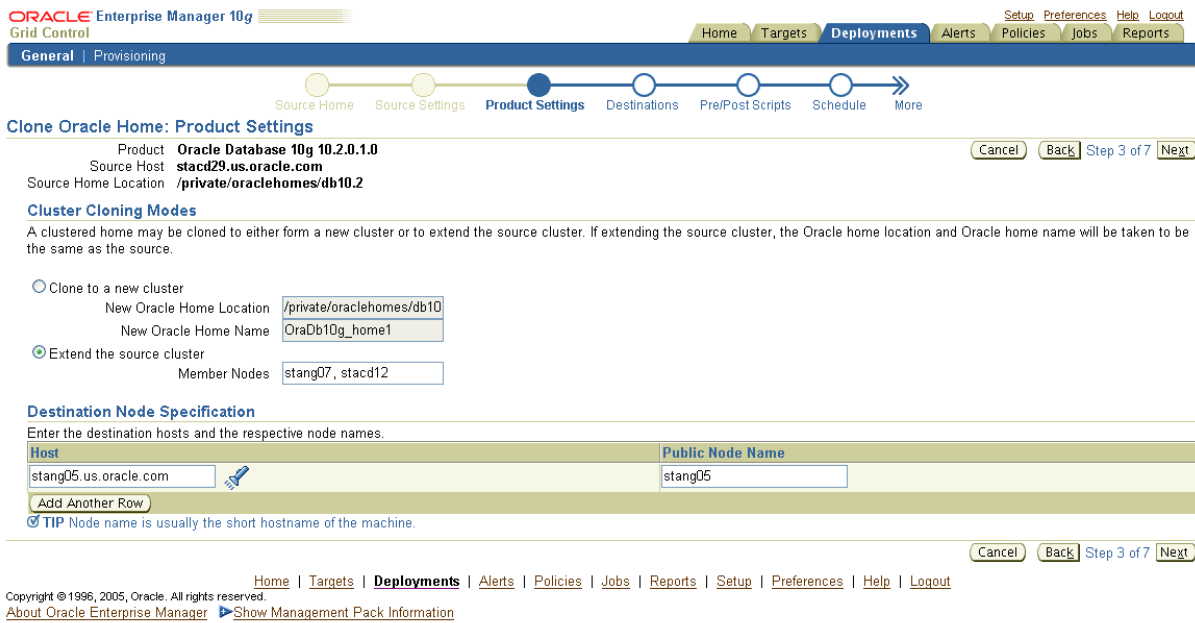
Grid Control allows administrators to define rules (or policies)—most of which are available out-of-box—that are responsible for governing the configuration of all the systems and software running across the enterprise. And because the natural configuration life cycle of any system means it will be changed—through installation of patches, adding files and directories, changing ports, editing its dependencies, and so on—once a system is rolled out, Grid Control continually audits it against its predefined policies. All changes are tracked so that administrators know when they are happening, who is making them, what changes are acceptable, and what changes must be corrected. This level of security and compliance through proactive auditing and enforcement is necessary to keep harmony in the continual flux that defines most data centers. It is also critical to maintaining order and meeting compliance standards as new systems are brought on line, existing systems are upgraded, or patches are applied.

## Error-Free Automated Provisioning: Linux, Oracle RAC, Certified Images, Patches, Upgrades

Systems administrators can spend up to 25% of their time installing and configuring new software. Using Grid Control's automated provisioning tools, businesses can reduce the time spent deploying new software by up to 90%. Grid Control allows administrators to store base images in a central library—pre-configured and

certified—from which new deployments can be based. Newly available with Release 2, administrators are able to deploy fully configured Linux systems from bare metal. These systems can be deployed with any desired software configuration, certified to the appropriate version and patch level—all with a few simple clicks. This capability is extraordinarily powerful for Oracle Real Application Clusters (RAC) and application cluster environments where adding capacity can be done quickly and without errors. Specific capabilities designed for Oracle RAC and cluster systems, minimizing the number of steps and complexities involved, are available with Release 2.

**Figure 2–2 Automated Oracle RAC Deployments and Extensions**



## Patch Management

Historically, managing the patches for any IT environment has been extremely tedious and difficult to control. In a Grid environment, where the number of systems and changes to those systems is widespread, manual patch management can become a full-time, often unproductive, job.

Using a direct link to Oracle, Grid Control proactively and regularly queries *My Oracle Support* for critical patches that have been released and notifies administrators through advisories of patches that should be applied. By correlating the available patches with the specific systems and software running on a particular site, administrators are notified of only applicable patches. When querying *My Oracle Support* directly, patches can be found either in the context of a specific target or, if desired, the administrator can query for a specific patch. Once the necessary patch is located, Grid Control can download and deploy it. Optionally, Grid Control can execute an end-user provided script to install the patch to all systems that require it. Each of these steps allows for quicker application of patches across the customer's enterprise. Grid Control extends this valuable patch management functionality to include support for OS patches—a critical component of systems maintenance and the biggest impediment preventing the installation of important software for system repair.

Figure 2–3 Searching for Patches in the Context of a Specific Target

ORACLE Enterprise Manager 10g  
Grid Control

Home | Targets | **Deployments** | Alerts | Setup | Preferences | Help | Logout  
Policies | Jobs | Reports

General | Provisioning

Select Patch | Select Destination | Set Credentials | Stage or Apply | Schedule | Summary

Patch: Select Patch Cancel Step 1 of 6 Next

To look up patches at Oracle MetaLink, enter search criteria and click Search. This may take a few moments depending on the number of matches found. From Search Results, select the patch to be applied and click Next.

Search by Criteria Search by Number

Product Family: Oracle Database  
Product: any  
Release: any  
(Find patches matching the release, and the most recent patchset for the release.)  
Patch Type: Patchsets  
Platform: Solaris Operating System (SPARC 32-bit)  
(If you run a 32-bit Oracle product on a 64-bit operating system, choose a 32-bit platform.)  
Language: Any  
Begin Date:   
End Date:   
(Example: 12/15/02)

Search Clear

Search Results

View Details View ReadMe Previous 10 121-129 of 129 Next

Select	Patch Number	Created On	Type	Product	Platform	Release	Interim Patch Applicable	Description	Status
<input checked="" type="radio"/>	4205997	Feb 25, 2005	Patchset	Workspace Manager	Generic Platform	10.1.0.3		WORKSPACE MANAGER RELEASE 10.1.0.5.0 PATCHSET	available
<input type="radio"/>	4205998	Feb 25, 2005	Patchset	Workspace Manager	Generic Platform	9.2.0.6		WORKSPACE MANAGER RELEASE 9.2.0.5.0 PATCHSET	available
<input type="radio"/>	4449379	Jun 29, 2005	Patchset	HTML DB	Generic Platform	10.1.0.3		PATCH SET FOR HTML DB (PATCH SET VERSION 1.6.1.00.03)	available
<input type="radio"/>	4469554	Jul 1, 2005	Patchset	Workspace Manager	Generic Platform	9.0.1.5		WORKSPACE MANAGER RELEASE 9.0.1.6.1 PATCHSET	available
<input type="radio"/>	4469558	Jul 1, 2005	Patchset	Workspace Manager	Generic Platform	9.2.0.5		WORKSPACE MANAGER RELEASE 9.2.0.5.1 PATCHSET	available
<input type="radio"/>	4469559	Jul 1, 2005	Patchset	Workspace Manager	Generic Platform	10.1.0.4		WORKSPACE MANAGER RELEASE 10.1.0.5.1 PATCHSET	available
<input type="radio"/>	765363	Jun 8, 2002	Patchset	Oracle Text	Solaris Operating System (SPARC 32-bit)	8.1.5.1		8.1.5.1:Oracle Context patchset for server version 8.1.5.1	available
<input type="radio"/>	953969	Jun 8, 2002	Patchset	Oracle Text	Solaris Operating System (SPARC 32-bit)	8.1.7.0		8.1.7.0:Intermedia Text Patchset 8.1.7.0.1	available
<input type="radio"/>	969660	Jun 8, 2002	Patchset	Oracle Text	Solaris Operating System (SPARC 32-bit)	8.1.6.0		8.1.6.0:Oracle interMedia Patch Set release 8.1.6.0.1	available

Cancel Step 1 of 6 Next

Home | Targets | **Deployments** | Alerts | Policies | Jobs | Reports | Setup | Preferences | Help | Logout

Using Grid Control to find, download, and deploy patches increases an administrator's efficiency by automating mundane, repetitive, day-to-day maintenance tasks as much as possible, allowing them more time for proactive systems work.

## Management for the Complete Oracle Grid Environment

While Enterprise Manager 10g provides the richest, most comprehensive management for Oracle products, it is also clear that there is no such thing as an Oracle-only environment. Customers have a wide variety of platforms, hardware, network, storage, and software components delivering critical functionality to the Oracle Grid, and they look for a single tool to monitor and manage this environment. With the latest release of Grid Control, there are two complementary initiatives that target extending the solution so customers can better leverage existing investments:

- Extending Grid Control's scope through native management extensions, and partner and custom plug-ins.
- Developing valuable integration solutions with existing management solutions to allow customers to leverage these investments without losing the unique value they gain from Grid Control.

## Extensions and Management Plug-ins

With Release 2, Grid Control dramatically broadens the reach of its management solution. New extensions and plug-ins enable customers to leverage their investment in Grid Control to manage all the pieces of the application solution in a vendor-neutral way—components such as firewalls, server load balancers, and non-Oracle databases and application servers. For instance, with Release 2, customers running Oracle WebLogic or IBM WebSphere will be able to seamlessly monitor and manage these system components with the similar, rich management functionality available for Oracle products.

Grid Control also includes extensions of the core Grid Control product for F5 server load balancers and NetApp Filers. Over the coming months, many more extensions will be generally released and available as easy-to-implement plug-ins to Release 2. These plug-ins will be both Oracle-developed (e.g., Checkpoint firewall, SQL Server, and others) and developed in partnership by third-party vendors, with support and validation by Oracle to provide a single console for primary monitoring and management of the Oracle Grid. In addition to the management plug-ins available out-of-box, Grid Control offers an open Command Line Interface and illustrative examples for easily adding new targets and customizing management solutions for customer-specific environments. This extended breadth of coverage gives customers the ability to use Grid Control for central management of their complete environment.

**Figure 2–4 Monitoring Oracle WebLogic Application Server Cluster**

The screenshot displays the Oracle Enterprise Manager 10g Grid Control interface for monitoring a BEA WebLogic Server Cluster. The main navigation bar includes Home, Targets, Deployments, Alerts, Policies, Jobs, and Reports. The current page is titled 'BEA WebLogic Server Cluster: stado65.us.oracle.com.medrecdomain.MedRecCluster' and shows data retrieved on Oct 9, 2005 at 3:40:41 AM PDT.

The interface is divided into several sections:

- Status:** Shows the cluster is 'Up' with 2 members, up since Oct 7, 2005 4:45:41 AM, and 100.00% availability. The member type is 'BEA WebLogic Managed Server'.
- Alerts:** A table showing the number of alerts by severity (Critical, Warning, Informational) for current and last 24 hours.
- Policy Violations:** A table showing the number of policy violations by severity (Critical, Warning, Informational) for current, last 24 hours cleared, and last 24 hours new.

Severity	Current	Last 24 hours
Critical (x)	0	0
Warning (y)	0	0
Informational (i)	0	0
Total	0	0

Severity	Current	Last 24 Hours		Distinct Rules Violated
		Cleared	New	
Critical (x)	0	0	0	0
Warning (y)	0	0	0	0
Informational (i)	0	0	0	0
Total	0	0	0	0

## Oracle Integrated Management Partnerships

In addition to extending the direct reach of Grid Control, bi-directional integration support for other leading management solutions is a key priority. Customers who have invested in non-Oracle management solutions derive increased value from Oracle's integration efforts. Direct integration between Grid Control Release 2 and products by HP, Micromuse, and EMC SMARTS are currently supported. As an example, the integration between HP OpenView Operations and Grid Control

provides both depth and breadth of management capabilities based on a common view of management data. This new level of integration allows customers to seamlessly leverage Oracle's strength in database, application server, and applications management and HP OpenView's strength in host and network management.

Other complementary partnerships designed to deliver unique capabilities to Oracle customers (including storage management and the ability to dynamically flex data center capacity) are underway and will be announced shortly.

## **Conclusion**

Grid computing is poised to revolutionize the economics of delivering applications and services across the enterprise. Enterprise Manager 10g allows customers to deliver heightened application performance and realize dramatic savings in IT administration and maintenance costs. Enterprise Manager's robust service level management functionality brings together business priorities with IT application performance requirements in ways it never has before. Rich configuration management and provisioning functionality eliminate the error-prone, manual tactics that plague most data centers and drive up the overall cost of management. In a time when administration resources are at a premium, and businesses are forced to tightly control IT spending budgets, Enterprise Manager 10g is a must-have tool for small and large enterprises that want to make the Grid a reality for their IT department.



---

---

## Overview of Enterprise Manager

Enterprise Manager enables users to manage their applications from top to bottom - from monitoring service levels to proactively isolating business exceptions before they escalate to addressing issues at any level in the IT stack. Users can manage the entire application lifecycle with comprehensive application quality management and compliance solutions.

This chapter describes the key features of Enterprise Manager:

- [System Monitoring](#)
- [Managing Groups](#)
- [Job System](#)
- [Information Publisher](#)
- [Compliance Management](#)
- [Extending Enterprise Manager](#)
- [Managing Targets](#)

### System Monitoring

Enterprise Manager provides comprehensive, flexible, easy-to-use monitoring functionality that supports the timely detection and notification of impending IT problems across your enterprise. For applications built on Oracle, Enterprise Manager offers the most comprehensive monitoring of the Oracle Grid environment - from Oracle Database instances to Oracle Real Application Clusters to Oracle Application Server Farms and Clusters.

For example, Enterprise Manager's monitoring functionality is tightly integrated with Oracle Database 10g manageability features such as server-generated alerts. These alerts are generated by the database itself about problems it has self-detected. Server-generated alerts can be managed from the Enterprise Manager console and include recommendations on how problems can be resolved. Performance problems such as poorly performing SQL and corresponding recommendations that are generated by the database's self-diagnostic engine, called Automatic Database Diagnostic Monitor (ADDM), are also captured and exposed through the Enterprise Manager console. This allows Enterprise Manager administrators to implement ADDM recommendations with ease and convenience.

Enterprise Manager comes with a comprehensive set of performance and health metrics that allow monitoring of key components in your environment such as applications, application servers, databases, as well as the back-end components on which they rely, such as hosts, operating systems and storage.

The Management Agent on each monitored host monitors the status, health, and performance of all managed components (also referred to as *targets*) on that host. If a target goes down, or if a performance metric crosses a warning or critical threshold, an alert is generated and sent to Enterprise Manager and to Enterprise Manager administrators who have registered interest in receiving such notifications.

Some metrics have predefined limiting parameters called thresholds that cause alerts to be generated when metric values exceed these thresholds. An alert indicates a potential problem indicating that a warning or critical threshold for a monitored metric has been crossed. An alert can also be generated for various availability states such as:

- Target is down.
- Oracle Management Agent monitoring the target is unreachable.

When an alert is generated, you can access details about the alert from the Enterprise Manager console. Administrators can be automatically notified when an alert is generated, and corrective actions can be automatically set up to resolve an alert condition.

*Corrective actions* allow you to specify automated responses to alerts. They ensure that routine responses are automatically executed, thereby saving administrators time and ensuring problems are dealt with before they noticeably impact users.

All monitoring information is accessed via the Enterprise Manager console, providing quick views into the health of your monitored environment.

The Enterprise Manager console home page shown in the figure below gives you an at-a-glance view of the overall status of your monitored environment. As shown in the following figure, the home page summarizes key monitoring areas such as availability across all managed targets, open alerts, policy violations, and recent problems with job executions. Links on this page allow you to drill down to detailed performance information.

The Resource Center is your central access point to Enterprise Manager documentation as well as the comprehensive technical resources of the Oracle Technology Network (OTN).





## Managing Groups

Today's IT operations staff are often responsible for managing a great number of components, such as databases, application servers, hosts, or other components. The Enterprise Manager Group Management System lets you combine components (called targets in Enterprise Manager) into logical sets, called *groups*. By organizing targets into groups, you can manage and monitor many components as one. The group management system enables you to organize, manage, and effectively monitor the potentially large number of targets in your enterprise.

The group management system allows you to:

- View a summary status of the targets within the group.
- Monitor outstanding alerts and policy violations for the group collectively, rather than individually.
- Monitor the overall performance of the group through performance charts.
- Perform administrative tasks, such as scheduling jobs for the entire group, or blacking out the group for maintenance periods.

## Job System

The Enterprise Manager Job System provides the capacity to automate routine administrative tasks and synchronize components in your environment so you can manage them more efficiently. A *job* is a unit of work that you define to automate commonly-run tasks. One of the advantages of jobs is that you can schedule a job to start immediately or start at a later date and time. You have the option to have the job run once or at a specific interval.

The Enterprise Manager Job System serves a dual purpose:

- Provides for the automation of many administrative tasks, for example, backup, cloning, and patching

- Allows users to create their own jobs using their own custom OS and SQL scripts
- Enterprise Manager provides predefined job tasks for database targets and deployments. A *job task* is used to contain predefined, unchangeable logic—for example, patch an application, back up a database, and so on. In addition to predefined job tasks, you can define your own job tasks by writing code to be included in OS and SQL scripts. In addition to submitting jobs to individual targets, you can submit jobs against a group of targets. Any job that you submit to a group is automatically extended to all its member targets and takes into account the membership of the group as it changes.

## Information Publisher

Information Publisher, Enterprise Manager's powerful reporting framework, makes information about your managed environment available to audiences across the enterprise. Reports are used to present a view of enterprise monitoring information for business intelligence purposes, but can also serve an administrative role by showing activity, resource utilization, and configuration of managed targets. IT managers can use reports to show availability of sets of managed systems. Executives can view reports on availability of applications (such as corporate email) over a period of time.

The reporting framework allows you to create and publish customized reports. Intuitive HTML-based reports can be published via the Web, stored, or e-mailed to selected recipients. Information Publisher comes with a comprehensive library of predefined reports that allow you to generate reports out-of-box without additional setup and configuration. The key benefits of using Information Publisher are:

- Provides a framework for creating content-rich, well-formatted HTML reports based on Management Repository data.
- Out-of-box reports let you start generating reports immediately without any system configuration or setup.
- Ability to schedule automatic generation of reports and store scheduled copies and/or e-mail them to intended audiences.
- Ability for Enterprise Manager administrators to share reports with the entire business community: executives, customers, and other Enterprise Manager administrators.

Information Publisher provides you with a feature-rich framework that is your central information source for your enterprise.

## Compliance Management

To have your enterprise run efficiently, it must adhere to standards that promote the best practices for security, configuration, and storage. Once these standards are developed, you can apply and test for these standards throughout your organization; that is, test for compliance. *Compliance* in this context is the conformance to standards, or requirements, or both.

Using Enterprise Manager, you can test the conformance of your targets for security standards, and configuration and storage requirements. By continually testing your systems, services, and targets, you are ensuring the best possible protection and performance for your system.

Enterprise Manager provides two types of compliance management: policies and policy groups. Together these define the optimal configurations of systems.

Policies and policy groups are similar in purpose, that is, they both provide rules against which managed entities are evaluated. However, there are differences:

- Rules within a policy group are managed as a set. They are viewed, maintained, evaluated, and reported, in the context of a policy group.
- Policies are evaluated as a set. Policies are viewed, maintained, and evaluated as standalone entities. Out-of-box policies and user-defined policies are available.

Whether you use the out-of-box policies and policy groups or customize policies and policy groups to meet your particular system requirements, any deviations to your systems or applications are reported. Examples of deviations include inappropriate settings and incorrect system configurations. Using the Information Publisher feature, you can view reports of any policy violations and policy group reports.

## Extending Enterprise Manager

Enterprise Manager allows you to monitor and manage a variety of components out-of-the-box.

Enterprise environments consist of a wide variety of components: OS platforms, hardware, software, network, and storage devices. All of these components work in concert to deliver critical information and functionality required to keep enterprise operations performing optimally and providing information to make important business decisions. While Oracle Enterprise Manager allows you to monitor and manage a variety of components out-of-the-box, you may want to monitor third party components or custom applications specific to your environment.

In addition, you can use the same mechanism used by Oracle and its partners to extend Enterprise Manager to monitor custom components via modular Management Plug-ins. The Management Plug-ins allow you to seamlessly monitor and manage non-Oracle components from the Grid Control console by providing an easy way to specify new classes of components for Enterprise Manager to monitor. Once you have created the new Management Plug-in, you can use the Enterprise Manager console to deploy the new plug-in throughout your enterprise.

Extending Enterprise Manager for monitoring additional components provides the following benefits:

- **Centralize management information in a single console:** When deployed, components defined by the Management Plug-in automatically appear in the Grid Control console. Being able to monitor all targets in your environment provides a consolidated view of your entire enterprise, thus allowing you to monitor and manage all components from a central point.
- **Extend Enterprise Manager's monitoring and management features to non-Oracle components:** Newly added components automatically inherit Enterprise Manager's powerful monitoring and management features, such as: alerts, policies, blackouts, monitoring templates, groups/systems, configuration management, and enterprise reporting.
- **Comprehensive Service-Level Management:** By managing all of your enterprise components with Enterprise Manager, you can more fully utilize the Service-Level Management features offered in Enterprise Manager.

**Tip:** For more details, see  
<http://www.oracle.com/technology/products/oem/extensions/index.html>

## Managing Targets

Enterprise Manager can monitor, administer, maintain, and manage different types of targets. For more details on these targets, refer to the following chapters:

- [Enterprise Configuration Management](#)
- [Service Management](#)
- [Database Management](#)
- [Change Management for Databases](#)
- [Middleware Management](#)
- [Oracle Collaboration Suite Management](#)
- [Host and Third-Party Target Management](#)
- [Oracle Beehive Management](#)
- [Virtualization Management](#)
- [Lifecycle Management](#)

---



---

## Enterprise Configuration Management

This chapter explains how Enterprise Manager simplifies the monitoring and management of the deployments in your enterprise, and contains the following section and subsections:

- [Hardware and Software Configurations](#)
- [Collected Configurations](#)
- [Viewing Configurations](#)
- [Comparing Configurations](#)
- [The Enterprise Configuration](#)
- [Client Configurations](#)
- [Accessing Configuration Pages in Grid Control](#)

### Hardware and Software Configurations

You can use Grid Control to view, save, track, compare, and search the configuration information stored in the Management Repository for individual hosts, databases, application servers, clients, and non-Oracle systems in the entire enterprise.

### Collected Configurations

Enterprise Manager collects configuration information for all hosts and the managed targets on those hosts that have a running Management Agent. The agent periodically sends the configuration information to the Management Repository over HTTP or HTTPS, allowing you to view up-to-date configuration information for your entire enterprise through Grid Control.

**Table 4–1** *Collected Configurations for Various Targets*

Target Type	Collected Configuration Information
Host <sup>1</sup>	<ul style="list-style-type: none"> <li>■ Hardware (includes memory, CPU, I/O device, and network information)</li> <li>■ Operating system (includes installed patches and patch sets)</li> <li>■ Oracle software (includes installed products and their components, patch sets, and interim patches applied using OPatch)</li> <li>■ Other software (includes all software registered with the operating system)</li> </ul>

**Table 4–1 (Cont.) Collected Configurations for Various Targets**

Target Type	Collected Configuration Information
Database <sup>2</sup>	<ul style="list-style-type: none"> <li>■ Database and instance properties</li> <li>■ Initialization and System Global Area parameters</li> <li>■ Tablespace, datafile, and control file information</li> <li>■ Redo logs, rollback segments, and high availability information</li> <li>■ Licensing information</li> </ul>
Application Server	<ul style="list-style-type: none"> <li>■ Installation type and version</li> <li>■ Repository information</li> <li>■ URLs (for Management Agent and Application Server Control)</li> <li>■ Component information (includes HTTP Server, OC4J, and Web Cache)</li> <li>■ OPMN information (includes ports, SSL information, and log levels)</li> </ul>
Client <sup>3</sup>	<ul style="list-style-type: none"> <li>■ Hardware</li> <li>■ Operating system (includes properties, file systems, patches)</li> <li>■ Software registered with the operating system</li> <li>■ Network data (includes latency and bandwidth to the Web server)</li> <li>■ Client-specific data that describes configuration for the browser used to access the client configuration collection applet</li> <li>■ Other client-oriented data items</li> </ul>
Non-Oracle Systems	<ul style="list-style-type: none"> <li>■ Hardware details including vendor, architecture, CPU, and I/O device information.</li> <li>■ Operating system details including name, version, software and package lists, kernel parameters, and file system information.</li> <li>■ OS Registered software including product name, vendor, location, and installation time.</li> </ul>
Enterprise <sup>4</sup>	Summary views for hardware, operating systems, Oracle Databases, Oracle Application Servers, and Oracle Collaboration Suite.

[Table 4–3](#) contains details on these summary views.

<sup>1</sup> The default collection period for host configuration information is 24 hours.

<sup>2</sup> The default collection period for database configuration information is 12 hours.

<sup>3</sup> Refer to "[Client Configurations](#)" in this chapter for more information.

<sup>4</sup> Refer to "[The Enterprise Configuration](#)" in this chapter for more information.

## Viewing Configurations

Using Grid Control, you can perform the following actions for targets such as hosts, databases, application servers, and clients:

- View the last collected and saved configuration
- Save configurations to a configuration file or to the Management Repository
- Search collected configuration data

- View the history of configuration changes
- Compare configurations (refer to "[Comparing Configurations](#)" in this chapter for more detailed information)

**See Also:** Grid Control online help for the pages described in [Table 4-2](#)

**Table 4-2 Viewing Configurations for Various Targets**

Target Type	Instructions for Viewing Configuration
Host	From the Home page for that host, click the <b>Configuration</b> subtab.
Database	From the Database Home page, click the <b>Maintenance</b> subtab, then click the <b>Last Collected Configuration</b> link under Software Deployments, Configuration.
Application Server	From the Application Server Home page, click the <b>Administration</b> subtab, then click the <b>Last Collected Configuration</b> link under Configuration.
Client	From the Deployments tab, click the <b>Client Configurations</b> link.
Enterprise	From the Deployments tab, select a summary view from the <b>View</b> list under Deployment Summary. Summary views are provided in <a href="#">Table 4-3</a> .

**Figure 4-1 Host Configuration Page**

**ORACLE Enterprise Manager 10g**  
Grid Control

Home | **Targets** | Deployments | Alerts | Setup | Preferences | Help | Logout

Collaboration Suites | Hosts | Databases | Application Servers | Web Applications | Services | Systems | Groups | **All Targets**

Host: [iinrnb05.us.oracle.com](#) Latest Data Collected From Target Oct 13, 2005 5:37:46 AM PDT [Refresh](#)

Home | Performance | **Targets** | Configuration

[Save](#) [History](#) [Compare Configuration](#) [Compare to Multiple Configurations\(Job\)](#)

**Hardware**

System Configuration	<b>i686</b>	Operating System	<b>Red Hat Enterprise Linux AS release 3 (Taroon Update 4) 2.4.21 27.EL.smp (32-bit)</b>
Hardware Provider	<b>Intel Based Hardware</b>	Packages	<b>1117</b>
Number of CPUs	<b>2</b>	Related Link	<a href="#">Operating System Details</a>
Memory Size (MB)	<b>5897</b>		
Related Link	<a href="#">Hardware Details</a>		

**Operating System**

**Oracle Software**

Product	Oracle Home	Installation Time
Oracle Application Server 10g 10.1.2.0.0	/home/sabburi/OraMid (OraMid)	Oct 5, 2005 12:28:59 AM
Oracle Application Server 10g 10.1.2.0.0	/home/sabburi/installs/oracluster (oracluster)	Oct 4, 2005 6:11:16 PM
Oracle Application Server 10g 10.1.2.0.2	/home/sabburi/midtier (oraclas_midtier)	Aug 31, 2005 1:35:11 AM
Oracle Application Server Integration B2B 10.1.2.0.2	/home/sabburi/OraMid (OraMid)	Oct 10, 2005 2:27:56 AM
Oracle Application Server Integration InterConnect 10.1.2.0.2	/home/sabburi/OraMid (OraMid)	Oct 5, 2005 3:39:52 AM
Oracle Application Server Integration InterConnect 10.1.2.0.2	/home/sabburi/midtier (oraclas_midtier)	Aug 31, 2005 3:13:42 AM
Oracle BPEL Process Manager 10.1.2.0.0	/home/sabburi/installs/oracluster (oracluster)	Oct 4, 2005 6:44:55 PM
Oracle Database 10g 10.1.0.2.0	/private/sabburi/databases/db10 (db_home)	Apr 19, 2005 6:32:32 PM
Oracle Management Agent 10.2.0.1.0	/scratch/sabburi/mm9_agent/agent10g (agent10g1)	Oct 7, 2005 6:03:50 AM

Related Link [Search Oracle Products Installed in Oracle Homes](#)

**OS-Registered Software**

Product	Vendor	Location	Installation Time
4Suite 0.11.1-14	Red Hat, Inc.	N/A	Mar 8, 2005 10:32:09 AM
a2ps 4.13b-28	Red Hat, Inc.	N/A	Mar 8, 2005 10:32:04 AM
acl 2.2.3-1	Red Hat, Inc.	N/A	Mar 8, 2005 10:31:18 AM
alchemist 1.0.27-1	Red Hat, Inc.	N/A	Mar 8, 2005 10:32:05 AM
alchemist-devel 1.0.27-1	Red Hat, Inc.	N/A	Mar 8, 2005 10:38:46 AM
am-utils 5.6.0.9-2	Red Hat, Inc.	N/A	Mar 8, 2005 10:36:33 AM
amanda 2.4.4p1-0.3E	Red Hat, Inc.	N/A	Mar 8, 2005 10:36:32 AM
amanda-client 2.4.4p1-0.3E	Red Hat, Inc.	N/A	Mar 8, 2005 10:36:35 AM
amanda-devel 2.4.4p1-0.3E	Red Hat, Inc.	N/A	Mar 8, 2005 10:38:51 AM
amanda-server 2.4.4p1-0.3E	Red Hat, Inc.	N/A	Mar 8, 2005 10:46:30 AM

Related Link [Search OS-Registered Software Installed on Hosts](#)

Home | Performance | **Targets** | Configuration

## Comparing Configurations

Grid Control gives you the tools to perform comparisons between configurations of the same target type. These comparisons are useful for quickly finding similarities and differences between two or more configurations.

You can compare:

- Two configurations in the Management Repository
- Two saved configuration files
- One configuration to multiple configurations
- A configuration in the Management Repository to a saved configuration file

When two target configurations are compared, all categories of collected configuration information are included. Grid Control presents the summary results of the comparison in a tabular format. More detailed information is available by drilling down from those summary results.

Comparisons between multiple configurations must be performed using the Enterprise Manager Job System.

**See Also:** "About Comparisons" in the Grid Control online help

## The Enterprise Configuration

Grid Control offers several summary views of your enterprise that you can use to monitor the state of your host and target configurations. "Enterprise" refers to the complete set of hosts and targets for which configuration information is stored in the Management Repository. You can also search enterprise configuration information using predefined or custom searches.

### Viewing the Enterprise Configuration

Table 4–3 lists the available enterprise configuration summary views.

**Table 4–3 Enterprise Configuration Summary Views**

Summary View	Description
Hardware	Shows the different types of host hardware and the number of hosts using each hardware type.  From the summary view, you can drill down to: <ul style="list-style-type: none"> <li>■ A list of hosts using each hardware type</li> <li>■ The operating system running on each host</li> </ul>
Operating Systems	Shows the different types of operating systems, the number of hosts running each type, and whether any operating system patches were applied.  From the summary view, you can drill down to: <ul style="list-style-type: none"> <li>■ A list of hosts running each operating system</li> <li>■ A list of the operating system patches, and whether or not they have been applied</li> </ul>



**Table 4–3 (Cont.) Enterprise Configuration Summary Views**

Summary View	Description
Database and Application Server Installations	<p>Shows the different versions of Oracle Databases or Application servers, respectively, the number of targets and installations for each version, and whether or not any interim patches were applied using OPatch.</p> <p>From the summary view for each database or application server version, you can drill down to:</p> <ul style="list-style-type: none"> <li>■ A list of the instances</li> <li>■ A list of the installations and Oracle home directories</li> <li>■ More information about the unique interim patches applied using OPatch</li> </ul>
Collaboration Suite Installations	<p>Shows the different versions of Oracle Collaboration Suite components, the number of targets and installations for each version, and whether or not any interim patches were applied using OPatch.</p> <p>From the summary view, you can drill down to more information about Oracle Collaboration Suite components.</p>

**See Also:** "Viewing the Enterprise Configuration" in the Grid Control online help

### Searching the Enterprise Configuration

In some cases, you may want to search your enterprise configuration to get answers to specific questions about your enterprise, such as:

- Which hosts have not had operating system patch 105181-05 installed?
- Which hosts have an Oracle version 9.0.1.0.0 database installed, and in what Oracle home directories are those databases installed?

Enterprise configuration searches query the enterprise configuration views in the Management Repository to find configuration information that satisfies the specified search criteria.

Enterprise Manager provides two types of enterprise configuration searches:

- **Predefined Searches**

Although these searches are predefined, you can modify the search criteria for each search, allowing you the flexibility to create specific search queries. Based on your search criteria, Grid Control creates the SQL query that searches the enterprise configuration views in the Management Repository.

Enterprise Manager provides the following predefined enterprise configuration searches:

- Search Oracle products, patch sets, and interim patches installed in Oracle homes
- Search software registered with the host operating system
- Search initialization parameter settings and setting changes
- Search tablespaces, datafiles, and recommended database settings
- Search database feature usage
- Search host operating system components, patches, property settings, and property changes

- Search host operating system and hardware summaries
- Search host file systems and network interface card configurations
- Search policy library
- **User-Defined Searches**

With a user-defined search, you specify the SQL query that searches the enterprise configuration views in the Management Repository. If you do not want to create the entire SQL query yourself, you can choose one of the predefined searches, make changes to the search criteria, then click **Search Using SQL** to display the SQL query that is executed. You can execute the query, view the results, then modify it and execute it again until it returns the desired results.

**See Also:** "Searching the Enterprise Configuration" in the Grid Control online help

### **Viewing the Enterprise Configuration History**

Enterprise Manager collects configuration information for a target and stores it in the Oracle Management Repository. Tracking historical changes to configurations is useful as it can help you diagnose problems.

Targets interact with each other and changes in one target may affect the behavior of other related targets. Tracking historical changes can be very important when trying to diagnose a problem as it provides clues as to which changes might have affected it. Hence, it is important to be able to see historical changes for a number of related targets.

The Enterprise Configuration History option allows you to see a merged history of a number of related targets for a specific period. For example, if the current value of a property seems to affect the value of another, the historical data helps you see when the value was last changed and what other values changed at that time.

## **Client Configurations**

A "client" represents an end-user or customer system—a system that is not part of your own IT infrastructure. A "client configuration" represents the configuration data collected about the end-user's system. These configurations differ from the internal deployments that you manage using Grid Control.

The Client System Analyzer (CSA) application allows Web server administrators to collect and analyze data from end-user systems. The client data is collected by an applet, diagnosed, and sent back to the CSA application. You can either use the CSA application that comes pre-installed with Grid Control, or you can deploy CSA independently to your Web server.

### **Client System Analyzer in Grid Control**

Using the pre-installed application allows you to collect client data without having to set up a separate Web server. The Management Agents collect, analyze, and upload the client data to the Management Repository. End users do not need login credentials to access Grid Control. Example usage scenarios include:

- End-users who call the Help Desk may be asked to navigate to the Out-Of-Box CSA page so that their system information is uploaded. The Technical Support Representative can then review the system information and offer solutions.
- The client's application can be changed to provide an "Upload my system information" link to the Client System Analyzer in the Grid Control application.

The link can specify certain configuration parameters, such as the URL to return to after the Client System Analyzer runs.

- The client's application can be modified to redirect its users to the Client System Analyzer in the Grid Control page during login or at other points in the application. Collected information can then be used from within Grid Control to obtain various bits of information about the client systems. Examples include most popular browser versions, or systems that do not have a necessary Operating System patch applied or do not have enough RAM.

### Client System Analyzer Deployed Independently

CSA can be deployed independently to any J2EE-capable Web server. This deployment strategy is appropriate when:

- Clients accessing CSA cannot reach or have limited access to a Grid Control deployment; for example, due to a firewall.
- Further customization to the CSA application is required, such as:
  - Custom rules can be supplied to the CSA application so that the end users have immediate feedback as to whether their systems satisfy certain constraints.
  - The behavior of the applet can be changed to collect additional information or to present end users with additional or different user interfaces.
  - The load on the Management Service Web servers needs to be reduced.

Both pre-installed and standalone types of deployments assign a configurable identifier called a Client Configuration Collection Tag to every client configuration collection. After the client configuration data has been collected by the client configuration collection applet and written to the Web server directory specified by the CSA application, you must configure Grid Control to collect the client configuration data and upload it to the Management Repository.

**See Also:** "Viewing a Client Configuration" in the Grid Control online help

"Configuring Enterprise Manager to Collect Client Configurations" in the Grid Control online help

## Accessing Configuration Pages in Grid Control

To access the configuration pages in Grid Control, click the **Deployments** tab, then click the links found under the Configuration section:

- Search
- Compare Configurations
- Compare to Multiple Configurations (Job)
- View Saved Configurations
- Import Configuration
- Host Configuration Collection Problems
- Refresh Host Configuration
- Configuration History

These links take you to pages where you can search, view, and compare configurations for your various targets.



---

---

# Database Management

This chapter introduces the concept of database management in the following sections:

- [Introduction to Database Management](#)
- [Monitoring Databases](#)
- [Administering Databases](#)
- [Masking Sensitive Data for Non-production Use](#)
- [Maintaining Databases](#)
- [Additional Maintenance Features](#)
- [Monitoring Oracle Real Application Clusters](#)

## Introduction to Database Management

Database management involves the monitoring, administration, and maintenance of the databases and database groups in your enterprise. Enterprise Manager is the premier tool for managing your database environment.

With Enterprise Manager, you receive:

- A complete set of integrated features for managing Oracle Databases
- Unparalleled scalability that lets you manage a single database or thousands of instances
- An intuitive management product that leads the industry in ease of deployment and use

## Database Control Versus Grid Control

Enterprise Manager provides two separate consoles that you can use to monitor your database: Database Control and Grid Control.

- Database Control is the Enterprise Manager Web-based application for managing Oracle Database 10g Release 2 (10.2). Database Control is installed and available with every Oracle Database 10g installation. From Database Control, you can monitor and administer a single Oracle Database instance or a clustered database.
- Grid Control is the Enterprise Manager console you use to centrally manage your entire Oracle environment. Within Grid Control, you access the multiple database targets using the **Targets** tab, then **Databases**.

## Database Home Page as the Management Hub

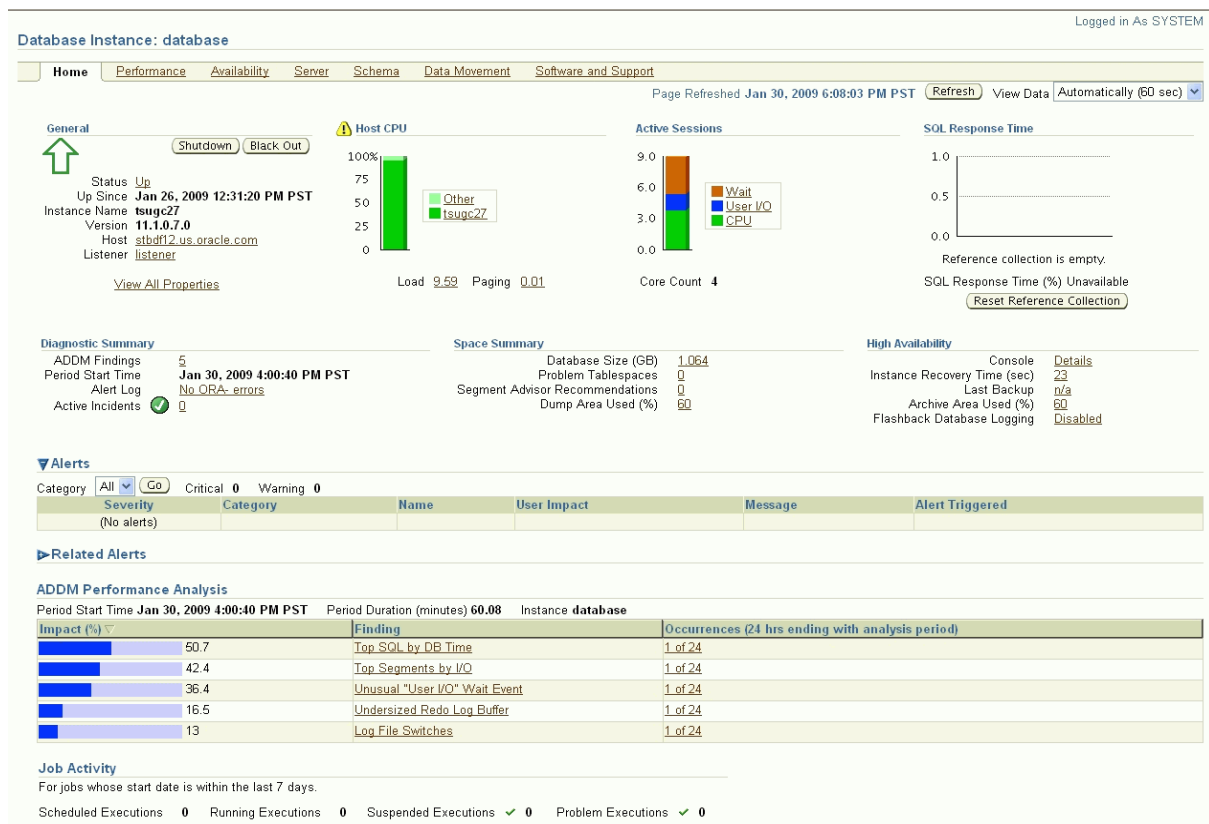
The Enterprise Manager Database Home page (Figure 5–1) shows important status and performance information about your database instance from a single source, including:

- A quick view of the status of the database and basic information about the database
- Relative CPU utilization of the Oracle host
- Amount of time the instance consumed using CPU and I/O, and the amount of time it consumed in bottlenecks
- Current response of the tracked set of SQL versus the reference collection response
- The number of Automatic Database Diagnostic Monitoring (ADDM) findings, the number of policy violations (Database Control only), and the Alert Log status
- Storage-related issues and recommendations for improved performance, and information about space violations and ADDM
- The most recent backup time and backup status
- Outstanding alerts and related alerts
- Rolled-up number of Oracle policy violations across all member targets in a system
- A quick view of database security
- Job executions that show the scheduled, running, suspended, and problem executions

Starting from the Database Home page, you can access additional details and management functions by drilling down through the user interface. In addition, the Database Home page provides a list of Related Links. From these links you can perform activities such as editing metric thresholds, analyzing job activity and metric collection errors, and accessing a number of advisors to help you improve the performance of your database.

**See Also:** "Oracle Database Home Page" in the Enterprise Manager online help

Figure 5–1 Database Home Page



## Monitoring Databases

Comprehensive database monitoring enables you to identify the problem areas in your database environment that are degrading performance. After you have identified the areas to improve, you can tune your databases' performance using the Enterprise Manager administration capabilities.

Enterprise Manager uses data from the Automatic Workload Repository (AWR) to display performance information and initiate database alerts. The user interface provides several real-time performance charts and drill-downs for the targets you manage. Both aggregate and instance-specific performance statistics are displayed using color-coded charts for easier viewing. To help you identify the source of a problem and resolve it, you can click a legend link next to a chart to display a detail page that provides comprehensive information.

The basic work flow in monitoring performance is to first go to the Database Performance page, which provides a high-level, comprehensive overview of important performance indicators.

## Assessing Database Performance

Database tuning is much easier and more effective when all the information for an in-depth performance diagnosis is available on a single screen. On the Database Performance page (Figure 5–2), which is accessible from the Database Home page, you can quickly view performance data for the instance and its underlying host. Additionally, you can analyze trends for all critical performance metrics and compare the trends to those of other database instances.

The most direct method for investigating and diagnosing performance problems from this page consists of accessing the Automatic Database Diagnostic Monitor (ADDM). After you start your Oracle database, the Automatic Workload Repository (AWR) begins taking snapshots of database activity once every hour by default, and ADDM executes when the snapshots are collected. ADDM uses these snapshots to perform a top-down analysis of your database activity and then provide tuning recommendations.

Another method for investigating and diagnosing performance problems consists of observing which of the wait classes beside the Average Active Sessions chart appears to be consuming an excessive amount of time, shown as a spike above the Maximum CPU line, then drilling down for more information. This provides a way for you to visualize the data for which ADDM has made its recommendations.

Either method assists you in diagnosing and resolving problems. The first automated method produces textual findings, whereas the second interactive method produces graphical findings.

Figure 5–2 Database Performance Page



The charts on the Database Performance page display current and recent metric information on a common time axis that enables visual correlation of metrics. These charts provide context-sensitive drill-downs where you can find additional details for rapid problem diagnosis.



**Table 5–1 Performance Page Charts**

Chart	Description
Host	This chart shows potential problems outside the database. The load average is a moving average of the run queue length. The run queue length indicates the level of contention for CPU time.
Average Active Sessions	This chart is the centerpiece of Oracle performance monitoring, and shows potential problems inside the database. It displays a profile of the amount of time sessions are either working or waiting to work in the database instance.  The Maximum CPU line on the chart helps determine how much of the CPU resource is being utilized. Session wait time accumulates above the CPU line, and the ratio of wait time to CPU time indicates how efficiently the system is working. Specific wait classes accumulating time indicate where you should focus your tuning efforts.
Instance Disk I/O	This chart shows the total requests that represent the rate at which the database instance is issuing read/write requests.
Instance Throughput	The charts show any contention that appears in the Average Active Sessions chart

The Additional Monitoring Links section enables you to access several related pages. Among those, the following pages have real-time diagnostic capabilities:

- Top Activity
- Top Consumers
- Instance Activity
- Historical SQL

**See Also:** "Performance Page" in the Enterprise Manager online help

## Diagnosing Problems Interactively

You can manually diagnose problems using an interactive method by first investigating a wait class that you think may be problematic, then drilling down from there for either SQL details or session details. You can alternatively investigate all wait classes in one location and then drill down from there. The following sections explain these methods.

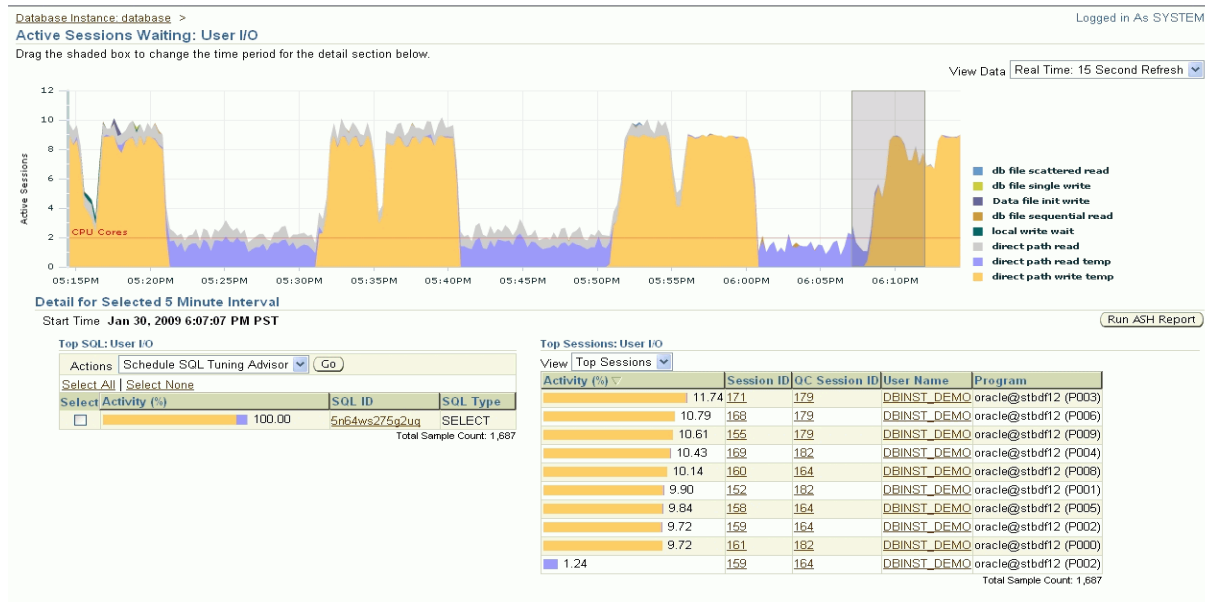
### Investigating a Wait Class

Wait class drill-downs enable you to investigate a specific wait class where it appears that time is accumulating. You can then determine whether the problem is caused by one or several SQL statements that are accumulating an unusual amount of time, or whether the problem is caused by one or several sessions.

The wait class drill-downs are called Active Sessions Waiting pages. The pages are sourced from the Active Session History (ASH), which samples session activity every second. It continuously records what is happening, such as which sessions are using CPU and which sessions are waiting on I/O.

Select a wait class link beside the Average Active Sessions chart to obtain detailed information for the wait class. For instance, if you click **User I/O**, the Active Sessions Waiting: User I/O page appears, as shown in [Figure 5–3](#).

**Figure 5–3 Active Sessions Waiting Page**



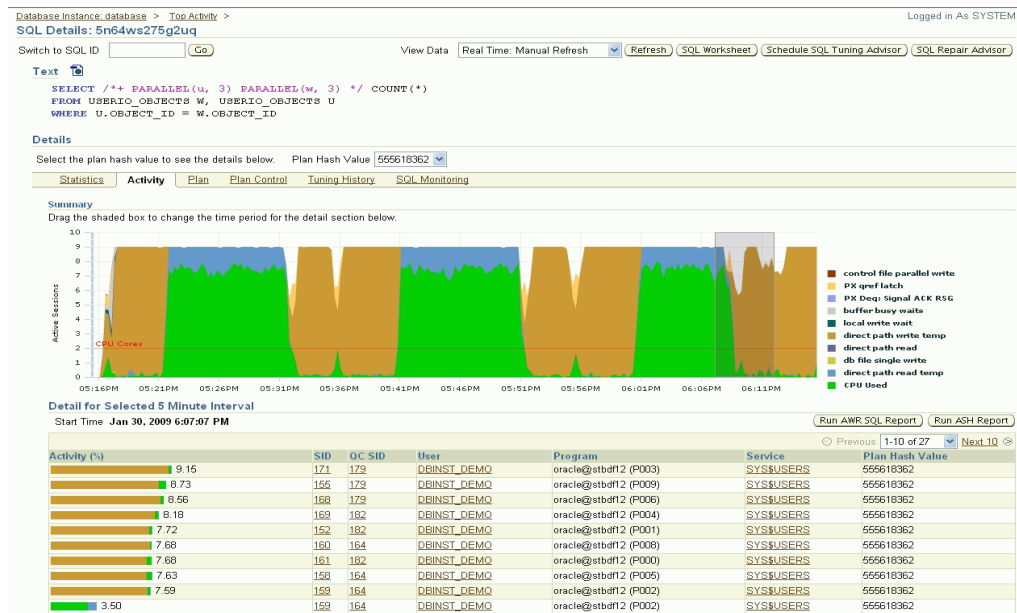
By default, the largest consumers of resources appear at the top of the detail tables. Look for skewed activity for either Top SQL or Top Sessions. If an excessive accumulation of activity appears to originate from an SQL source, you can click its associated SQL ID to go to the SQL Details page, which shows the SQL statement and activity. If excessive accumulation appears to originate from a session source, you can click its associated Session ID to go to the Session Details page, where you can kill the session if needed. You can also view the wait events associated with the session.

**See Also:** "Active Sessions Waiting Page" in the Enterprise Manager online help

**Viewing SQL Details** The SQL Details page, shown in Figure 5–4, displays the SQL statement you select in the Active Sessions Waiting page. You can also:

- Examine the activity of this SQL statement over time.
- View SQL-level statistics.
- Investigate the SQL plan.
- Access any prior Tuning Advisor run.
- Schedule a Tuning Advisor run.

Figure 5–4 SQL Details Page

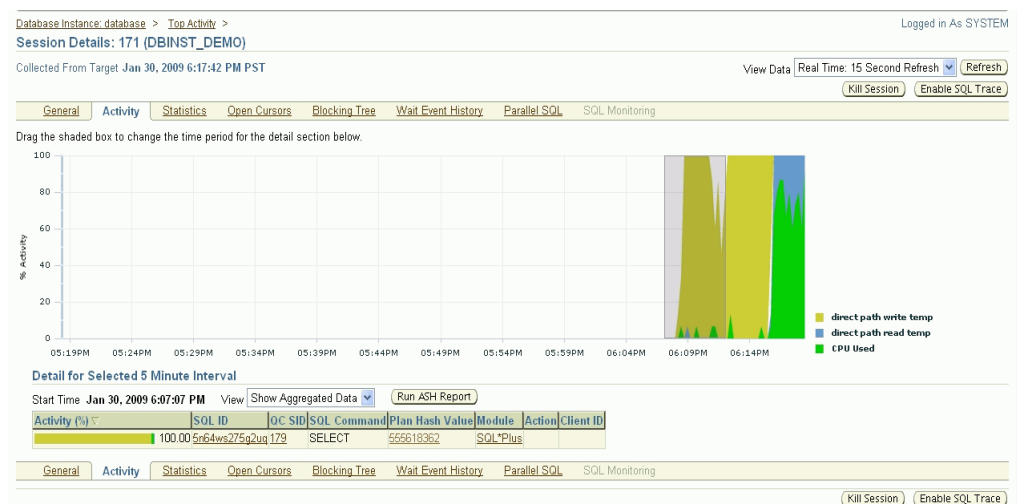


**See Also:** "SQL Details Page" in the Enterprise Manager online help

**Viewing Session Details** The Session Details page, shown in Figure 5–5, displays the wait events associated with the session you selected in the Active Sessions Waiting page. You can also:

- View the current values for metrics associated with the current session.
- View the list of currently open cursors in the selected session, including their hash value and SQL text.
- View sessions that are blocking other sessions.

Figure 5–5 Session Details Page

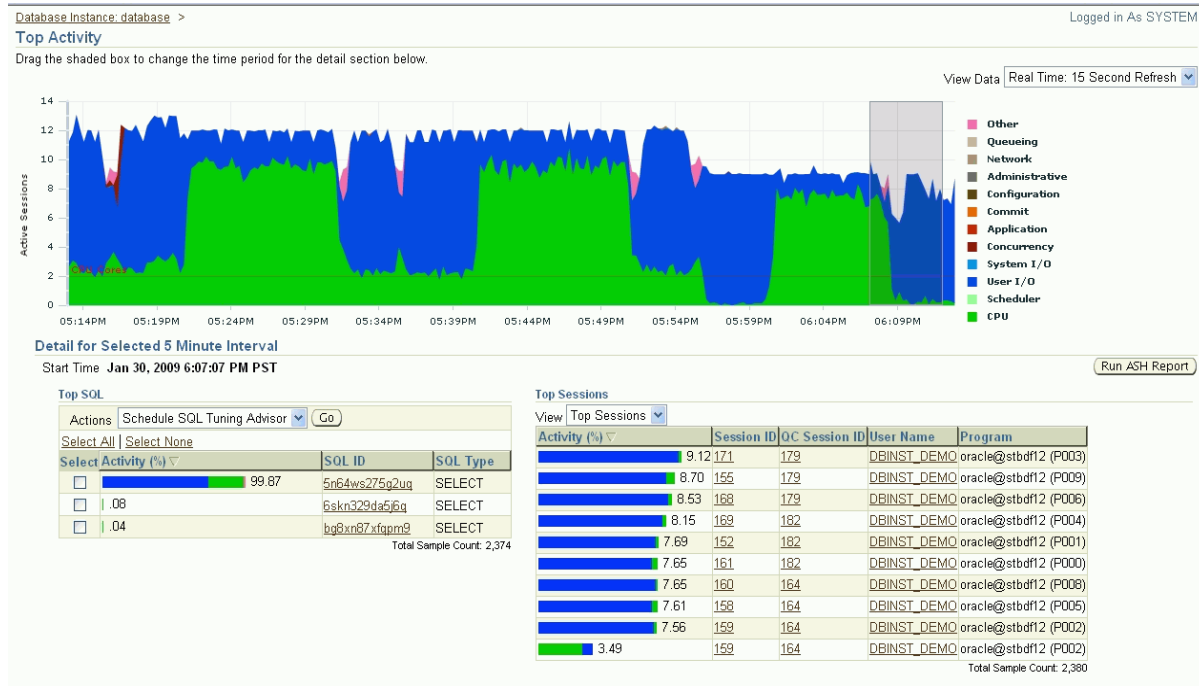


**See Also:** "Session Details Page" in the Enterprise Manager online help

## Viewing Top Activity

The Top Activity page, shown in [Figure 5–6](#), is essentially the sum of all wait class drill-downs. The interface format is the same as the Active Sessions Waiting page, but rather than showing the average active sessions for a particular wait class, the Top Activity page shows the average active sessions for all wait classes. The page also shows the Top SQL and Top Sessions consumers for all wait classes. As with the Active Sessions Waiting pages, look for skewed activity for either Top SQL or Top Sessions.

**Figure 5–6 Top Activity Page**



**See Also:** "Top Activity Page" in the Enterprise Manager online help

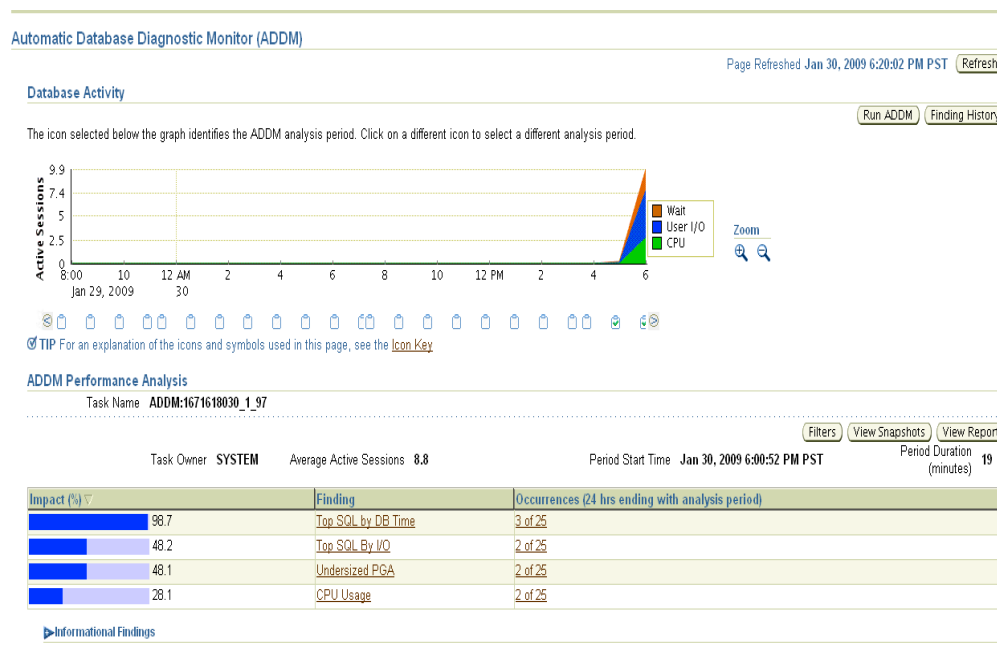
## Diagnosing Problems Automatically

As mentioned in [Assessing Database Performance](#), you can use the Automatic Database Diagnostic Monitor (ADDM) to automatically investigate and diagnose performance problems from either the Database Home page or the Database Performance page. Using regularly scheduled snapshots of the database activity, ADDM identifies the most resource-intensive components or operations, and determines whether or not these components or operations are acting as performance bottlenecks. If one or more problems have occurred, ADDM diagnoses these potential problems and provides advice, which may recommend that you run an advisor or change your database configuration.

You can view ADDM findings over a predetermined period of recent time or at the present time. You can view ADDM findings for the most recent time interval by clicking the **ADDM Findings** link in the Diagnostic Summary section of the Database Home page. You can also view ADDM findings for this interval by clicking the small ADDM icon below the Average Active Sessions chart in the Database Performance page. Either option displays the Automatic Database Diagnostic Monitor page, shown in [Figure 5–7](#). To determine ADDM findings at the present time, you can click the **Run ADDM** button on the Database Performance page.

By default, the database takes snapshots at 60-minute intervals. You can use the Automatic Workload Repository to change the snapshot interval, ranging from 10 minutes to 2 hours.

**Figure 5–7 Automatic Database Diagnostic Monitor Page**



**See Also:** "Performance Finding Details Page" in the Enterprise Manager online help

## Using Additional Diagnostic Pages

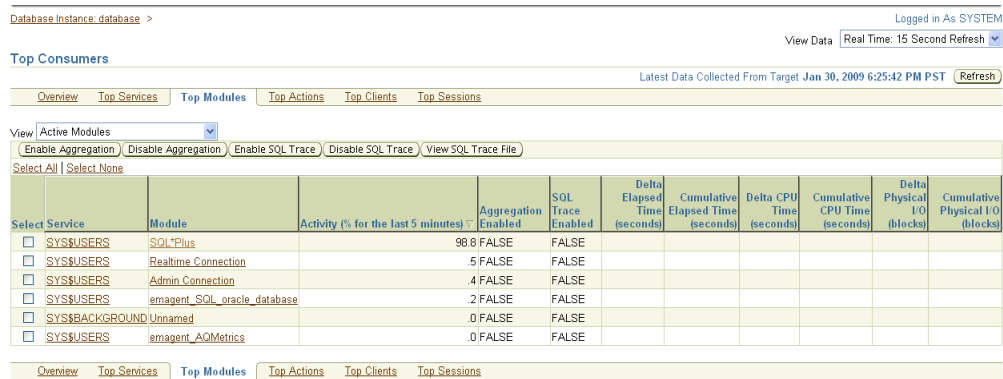
Besides the primary diagnostic pages discussed above, there are other important secondary pages that can assist you in diagnosing and correcting performance problems. The following sections discuss these diagnostic pages:

- Top Consumers
- Instance Activity
- Historical SQL

### Top Consumers

The Top Consumers link provides global summary information for the top database consumers of system resources. You can access detailed metrics data for a specific top consumer, such as sessions, services, modules, and clients. This enables you to pinpoint the most problematic areas on which to focus your database tuning efforts. [Figure 5–8](#) shows the Top Modules page of Top Consumers, where you can perform tasks such as enabling and disabling aggregation and SQL traces for the consumer. You can use an SQL trace to trace statistics of SQL statements, such as CPU time, elapsed time, and the Explain Plan.

**Figure 5–8 Top Consumers Top Modules Page**

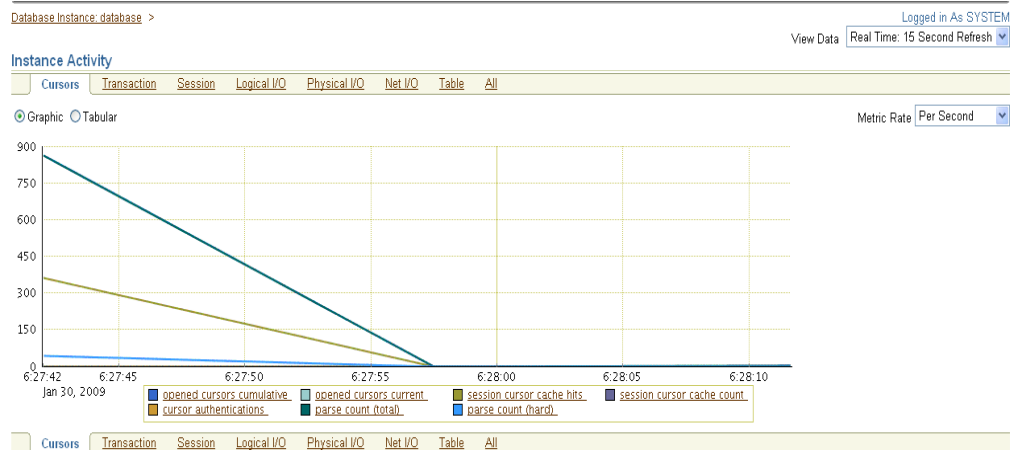


**See Also:** "Top Consumers Page" in the Enterprise Manager online help

### Instance Activity

The Instance Activity link displays database activity for specific data about groups of metrics, such as cursors, sessions, and transactions (Figure 5–9). For instance, the Cursors metric group displays information about opened and session cursors along with authentications and parse counts. The legend links below the chart in the graphic view, or the name links in the table view, enable you to access the Top Sessions page for more detailed information.

**Figure 5–9 Instance Activity Page**



**See Also:** "Instance Activity Page" in the Enterprise Manager online help

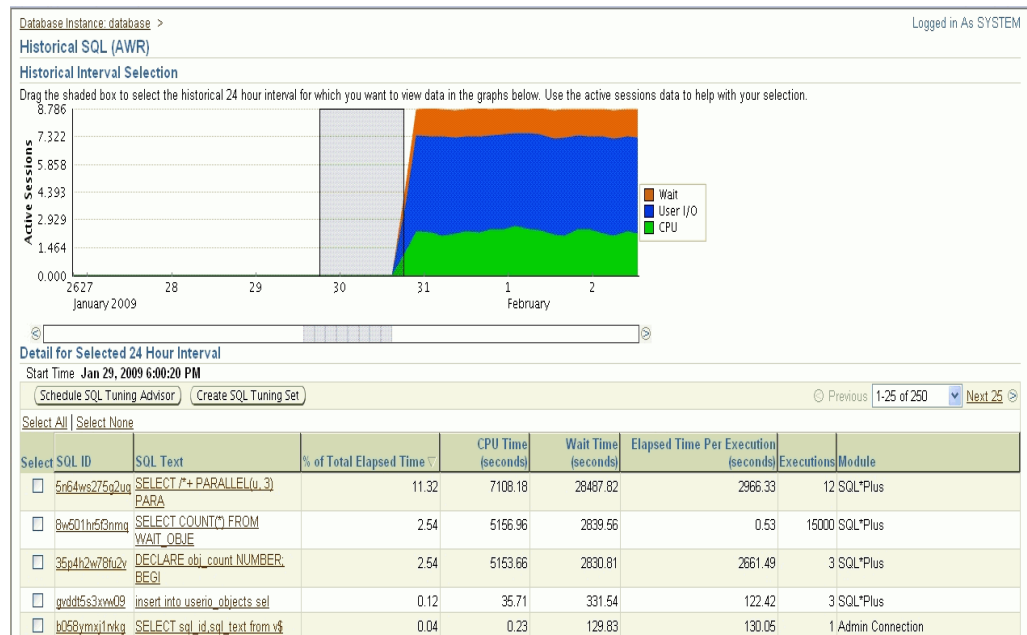
### Historical SQL

The Historical SQL (AWR) link displays statements stored in the Automatic Workload Repository (AWR) for a 24-hour period. This link is available in the Additional Monitoring Links section when you select the Historical data view (Figure 5–10).

The table at the bottom of the page shows an analysis of all SQL statements for performance and resource consumption. You can select a statement link to see SQL details (statistics, activity, SQL plan, and tuning information) for the statement. You

can also run the SQL Tuning Advisor to receive recommendations for one or more statements.

**Figure 5–10 Historical SQL Page**



**See Also:** "Historical SQL (AWR) Page" in the Enterprise Manager online help

## Using Other Performance Tools

Enterprise Manager offers several tools that provide additional assistance ranging from viewing the relationships between system components to collecting system statistics when database performance is degraded. The following sections discuss these tools:

- Topology Viewer
- Metric Baselines
- Memory Access Mode
- Hang Analysis

### Topology Viewer

Enterprise Manager provides a Topology Viewer for several applications. The Topology Viewer allows you to view the relationships between components, nodes, or objects within different Oracle applications. You can zoom, pan, see selection details and summary information, and evaluate aggregate components. Individually distinct icons are used for each object type, and standardized visual indicators are used across all applications.

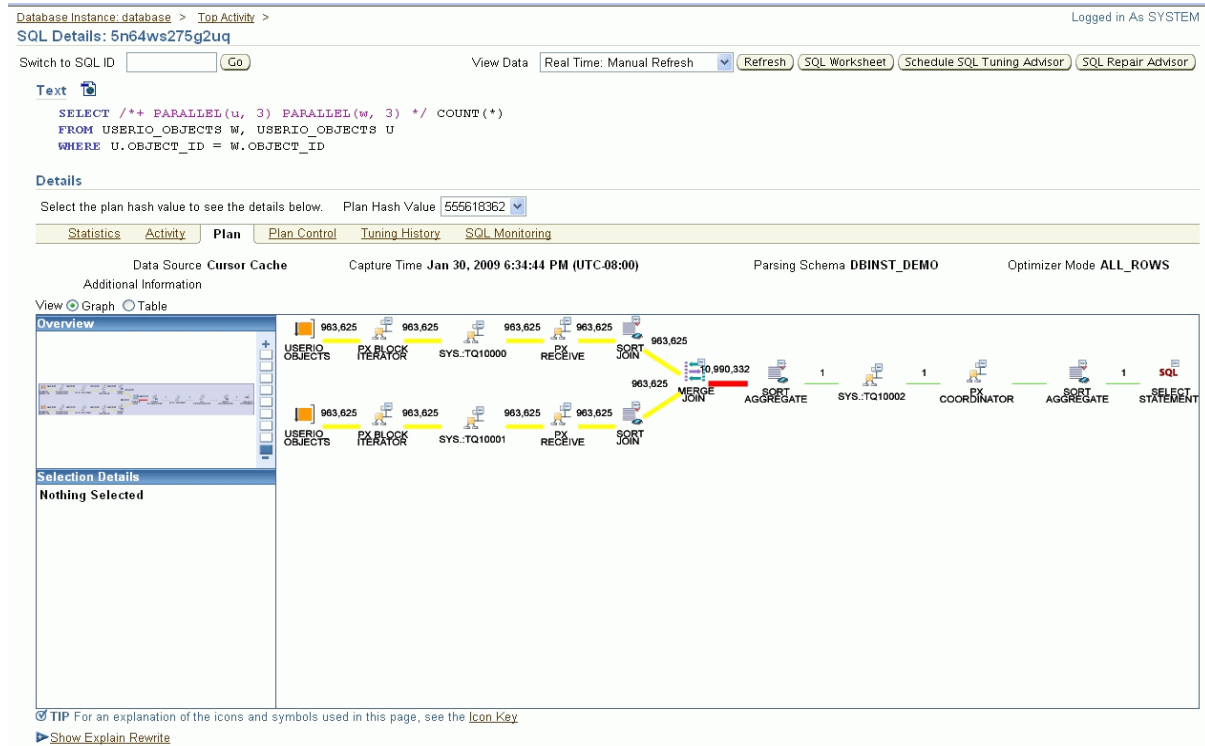
The Topology Viewer is available for the following database applications:

- Scheduler
- SQL Details
- SQL Explain Plans

- Oracle Real Application Clusters
- Cluster Databases

Figure 5–11 shows the Topology Viewer for the SQL Details Plan page. The Topology Viewer provides a graphical representation of your SQL steps as modeled in Enterprise Manager.

**Figure 5–11 Topology for SQL Details Plan**



### Metric Baselines

From the Database Performance page, you can set a metric baseline—a time period associated with a target that you can use as a reference for evaluating target performance (Figure 5–12). Statistics are computed over the baseline period for specific target metrics. You can use these statistics to automatically set metric thresholds for alerts, as well as to normalize graphical displays of system performance.



**Figure 5–12 Metric Baselines Page**

**Metric Baselines**

Accumulated trailing data **16 days** (of 91 day AWR retention)  
[\(Change AWR Retention\)](#) [Disable Metric Baselines] [Register Metrics]

**Select Active Baseline**

No active baseline

Moving window baseline Trailing 21 days ▼ Time Group: By Day and Night, over Weekdays and Weekend ▼

TIP Moving baselines (recommended) accommodate evolving systems.

Static metric baseline ▼ Time Group: N/A

Time Period: N/A

[Set Adaptive Thresholds] [Revert] [Apply]

**See Also:** "Metric Baselines Page" in the Enterprise Manager online help

### Memory Access Mode

When the database server experiences performance issues, diagnostic queries could further impact the system. A new feature, called Memory Access Mode, assists you in diagnosing performance-related problems by collecting system statistics even when the database is either slow or hung.

Instead of relying on the normal SQL engine, the data is retrieved by accessing tables directly from the Shared Global Area (SGA). The data collection is quick and does not further impact a system that is already slow. Potentially, this allows for sub-second sampling of performance metrics without causing any additional load on a system.

You can disable the standard SQL Access Mode and switch to Memory Access Mode by clicking the **Monitor in Memory Access Mode** link in the Related Links section.

**See Also:** "Database Performance Page" in the Enterprise Manager online help

### Hang Analysis

A new feature, called Hang Analysis, enables you to diagnose locking issues that might either slow down a system or cause hanging. Normally, if a system is slow or in a hung state, the diagnostic queries are also either extremely slow or never return a result. This utility bypasses the typical query engine instead, and leverages the Oraclebug API to return results quickly even on systems that might be seemingly in a hung state.

The Hang Analysis page shows a visual map of the sessions that are either blocking or are blocked. A tree view of the sessions is displayed, and the problem session that is blocking other sessions is located at the root of the tree. Each session is shown color-coded, signifying how long the session has been blocked. Clicking on the session box brings up another page that shows session details. Using information from this page, you can cancel the problem session and return the system to its normal state.

You can access the Hang Analysis page by clicking the Hang Analysis link in the Additional Monitoring Links section.

**See Also:** "Hang Analysis Page" in the Enterprise Manager online help

## Administering Databases

Oracle Enterprise Manager effectively keeps your Oracle Databases available and running efficiently. Enterprise Manager can help database administrators perform everyday tasks. Specifically, it provides a graphical user interface for managing database storage structures and schemas. As with database monitoring, administration of the Oracle Database begins with the Oracle Database Home page. From this page, you can display an overview of your database properties and performance. However, you can also use the Administration section of the page to perform common administration tasks such as the following:

- Allocating system storage and planning future storage requirements for the database system
- Creating and managing primary database storage structures (tablespaces)
- Creating and managing primary objects (tables, views, indexes)
- Enrolling users and maintaining system security by controlling and monitoring user access to the database
- Backing up and restoring the database

Just as Enterprise Manager monitoring identifies problem areas in your database and database groups, you can administer your database using the Enterprise Manager administration tools. The administration tools allow you to manage database objects and initiate database operations inside an Oracle Database. The following sections provide an overview of the database administrative features available to you in Enterprise Manager.

### Managing Storage Objects

As an administrator, you can use the administration tools in Oracle Enterprise Manager to optimize database performance. Using these tools, you can manage storage structures such as control files, tablespaces, datafiles, and archive logs. In addition to viewing, editing, and deleting these structures, you can also perform other functions such as making tablespaces locally managed, displaying the dependencies of a datafile, or backing up a control file to a trace.

### Using Database Configuration Features

Oracle Enterprise Manager incorporates many features that help you to manage your database configuration. For example, you can use Enterprise Manager to manage the memory size of the system global area and program global area of your system. You can also use the Undo Management feature to provide a means of explicitly specifying the amount of undo information to retain and ensure that undo information is not overwritten.

You can create or edit initialization parameters for the current database, setting these parameters to specific values to initialize many of the memory and process settings of an Oracle Database instance. You can also display a list of database features, showing how often the features are used in the operation of the database. Usage information can be utilized by support groups and other organizations to gain knowledge about how the system is being used, and to help apportion resources as necessary.

### Using Automatic Storage Management

Enterprise Manager uses Automatic Storage Management (ASM) to automate and simplify the layout of datafiles, control files, and log files. Automatic Storage

Management (ASM) simplifies database administration by managing disk groups instead of thousands of individual database files using disk groups. Database files are automatically distributed across all available disks, and database storage is rebalanced whenever the storage configuration changes. ASM internally creates, deletes and manages files as and when needed using the storage space from the disk groups. ASM also includes storage reliability features such as mirroring and parity protection commonly found in Logical Volume Managers. These features are available to an Oracle database or to Real Application Clusters (RAC).

An ASM instance is a special Oracle instance for coordinating ASM activities. An instance provides services to an Oracle database within a node. On a single node, there is typically a single ASM instance on the node which manages all disk groups. In a RAC environment, there is typically one ASM instance on each node which manages all disk groups for its node in a coordinated manner.

Automatic Storage Management files can be configured for use as general storage containers or volume devices. Oracle File System is a general-purpose file system that supports all customer application data files, both database and non-database including those associated with an Oracle home. This support is on single host and cluster configurations.

## Converting Single Instances to Oracle Real Application Clusters

Oracle Real Application Clusters (RAC) provides a high-availability database environment spanning multiple hosts. Each cluster can be comprised of multiple cluster databases, each of which consists of multiple cluster database instances. A cluster database is available as long as one of its instances is available. You can use Enterprise Manager to asynchronously convert a single-instance database to an Oracle RAC database.

## Converting to a Locally-Managed Tablespace

An added feature allows you to convert a dictionary-managed tablespace to a locally-managed tablespace, thereby simplifying space allocation, improving performance, and reducing reliance on the data dictionary.

## Controlling Resources with Resource Manager

The Database Resource Manager controls the distribution of resources among various sessions by controlling the execution schedule inside the database. By controlling which sessions run and for how long, the Database Resource Manager can ensure that resource distribution matches the plan directive and hence, the business objectives. You can use the Database Resource Manager to automatically assign sessions to specific consumer groups by setting up mappings between session attributes and consumer groups. You can also map consumer groups in specified categories to users, client programs, modules, or services.

Resource consumer groups let you group user sessions together by resource requirements. Resource consumer groups are different from user roles; one database user can have different sessions assigned to different resource consumer groups. You can then use a resource plan to specify how the resources are to be distributed among various resource consumer groups.

Resource plans specify the resource consumer groups belonging to the plan and contain directives for how resources are to be allocated among these groups. Plan information is stored in tables in the data dictionary. Several views are available for

viewing plan data. In addition to containing resource consumer groups, a plan can contain subplans. Use Enterprise Manager to manage all aspects of the resource plan.

## Tracking Statistics to Improve Database Performance

The Workload Repository provides you with a mechanism for gathering database statistics for specific time intervals. You can use the Optimizer Statistics feature of Enterprise Manager to simplify the management of Optimizer Statistics operations such as gathering, restoring, deleting, locking, and unlocking statistics. Use these statistics to improve the performance of your SQL statements.

## Using Oracle Scheduler

Oracle Enterprise Manager uses the Scheduler to enable database administrators and application developers to control when and where various tasks occur in the database environment. Using the Scheduler helps to improve the management and planning of these tasks. By separating a task into its component parts such as time, location, and database object, the Scheduler provides an easier way to manage your database environment. Database administrators can schedule and monitor recurring database maintenance jobs, such as backups or nightly data warehousing loads and extracts, and they can schedule job execution based on time or events.

With Enterprise Manager, you can enable and disable Scheduler jobs, alter the settings for existing jobs, start or stop current jobs, and view Scheduler information.

## Working With Database Schemas

A schema is a collection of database objects comprised of logical structures that directly refer to the data in the database. Schema objects include structures such as tables, views, and indexes. These schema objects can be created and managed using the tools available in Oracle Enterprise Manager.

## Managing Database Objects

Oracle Enterprise Manager provides a comprehensive set of tools that allows you to manage all aspects of database directory objects such as tables, indexes, and views. You can use the tools available in Enterprise Manager for fundamental tasks such as creating, editing, and viewing object properties, but you can also use Enterprise Manager for more comprehensive tasks such as running the Segment Advisor to evaluate a table for block and space usage and to determine whether space can be saved by shrinking highly fragmented segments. The space gained by the implementation of these recommendations is returned to the table.

Indexes are optional structures associated with tables that can be created to increase the performance of data retrieval. When you manage indexes in Enterprise Manager, you can perform functions such as shrinking segments to compact segments and free the recovered space to the current tablespace. Alternatively, you can also eliminate space problems by reorganizing space usage while changing the storage settings and location of the index.

Views are customized presentations of data in one or more tables or other views. In addition to creating, deleting, and managing views, you can also view the objects that are dependencies of the view. The Dependencies table shows you the Object Name and the Object Type that is dependent on the current view. Conversely, Enterprise Manager can also show you the objects on which the current view is dependent.

In general, use the Action menu on the Search page or Property page of any database object to see a list of the available functions you can perform on that object.

Enterprise Manager allows you to manage program structures, such as packages, package bodies, functions, and triggers in the same way. In addition to creating and viewing these elements, you can compile Java Sources, for example, or generate the data definition language (DDL) code for a specified object such as package. For more information about the various functions available for a specific object, you can refer to the online Help.

## XML Database Features in Oracle Enterprise Manager

XML, eXtensible Markup Language, is the standard way to identify and describe data on the Web. Oracle XML DB treats XML as a native datatype in the database. Oracle XML DB offers a number of easy ways to create XML documents from relational tables. The result of any SQL query can be automatically converted into an XML document. Oracle also includes a set of utilities, available in Java and C++, to simplify the task of creating XML documents.

## Managing Users and Privileges

Oracle includes security features that control how a database is accessed and used. Privileges and roles control user access to data and the types of SQL statements that can be executed. There are three types of privileges and roles: system privileges, object privileges, and roles. The system privilege is an Oracle-defined privilege usually granted only to and by administrators. It allows users to perform specific database operations. An object privilege is an Oracle-defined privilege that controls access to a specific object. A role can be created by users, usually administrators, and is used to group together privileges and other roles. This facilitates the granting of multiple privileges and roles to users.

Privileges and roles can be granted to other users by users who possess the power to do so. Oracle Enterprise Manager allows you to create and manage users, roles and profiles. You can also expire passwords and lock or unlock users by applying those actions against one or more users. When managing roles, you can display all users and roles assigned to the specified role by using the Show Grantees function.

Auditing is the monitoring and recording of selected user database actions. It can be based on individual actions, such as the type of SQL statement run, or on combinations of factors that can include name, application, time, and so on. Security policies can cause auditing when specified elements in an Oracle Database are accessed or altered, including content. You can set up and adjust audit settings easily within the Enterprise Manager interface. With Enterprise Manager, you can view the database audit configuration and manage audited objects, privileges, and statements. You can also view the content of the audit trail. Out-of-box, Enterprise Manager also provides audit reports to help you monitor successful and failed login attempts, as well as SYSUSER operations.

## Managing Materialized Views

Materialized views are schema objects that can be used to summarize, compute, replicate, and distribute data. They are suitable in various computing environments such as data warehousing, decision support, and distributed or mobile computing. Enterprise Manager lets you create and manage materialized views and provides a set of additional tools that allows you to perform specific actions against the view. The Explain Materialized View feature helps you to determine what is possible with a materialized view by indicating whether a materialized view is fast-refreshable, what

types of query rewrites you can perform with this materialized view, and whether PCT refresh is possible.

Oracle maintains the data in materialized views by refreshing them after changes are made to their master tables. The refresh method can be incremental (fast refresh) or complete. For materialized views that use the fast refresh method, a materialized view log keeps a record of changes to the master tables. You can use Enterprise Manager to advise you when a query can be rewritten. You can then take the appropriate action required to make a query rewrite. Query rewrite transforms a SQL statement expressed in terms of tables or views into a statement accessing one or more materialized views that are defined on the detail tables.

Materialized views can be refreshed either on demand or at regular time intervals. You can use the Action menu in Enterprise Manager to manually refresh a materialized view.

## About Change Management

A dictionary baseline is an object containing a set of database definitions captured at a certain time. Baselines are stored internally in the Enterprise Manager repository, and are in a form that other Change Management applications can use. You can use Enterprise Manager to capture the database object definitions at a particular point in time and create reusable baseline scope specifications. By capturing a dictionary baseline, you can compare different database objects at different points in time and track changes applied to the database objects.

A dictionary comparison identifies differences in database object definitions between two databases, a database and a baseline, or two schemas within a single database/baseline. With Enterprise Manager, you can compare, view, and track differences between two sets of database object definitions at different points in time. The comparison result displays identical objects, objects/attributes that are different, and objects present either in the left or right source.

A dictionary synchronization synchronizes differences in database object definitions between two databases or a database and a baseline.

Synchronizations are generated using synchronization specifications. For synchronizations, the scope specification does not include the names of individual objects. It can only specify the types, and the schemas to be included or excluded. You can additionally supply a prefix to limit the objects selected to those whose names start with that prefix.

Dictionary synchronizations synchronize differences in any attribute value between objects of any type. Use synchronization specifications to create multiple versions of a synchronization. Each version has a unique version number and a synchronization date. Use these versions to associate synchronizations of database/schemas made over time.

For more information about these features, see [Change Management for Databases](#).

## Using Advisors in Oracle Enterprise Manager

Advisors are procedures that you can invoke, or Enterprise Manager can invoke internally, that designate a specific object for analysis. An advisor can report on a variety of aspects of the object, and describe a recommended action for each condition worthy of user intervention. The advisor might report that the condition can be corrected by an automated task that it provides.

Some advisors also provide what-if analysis for specific situations. For example, the Undo Advisor provides an analysis of the impact of changing the retention period for undo records on the size of the undo tablespace. Additionally, the Memory Advisor graphically displays the impact on performance of changing the size of a component of the SGA.

You can invoke an advisor from the Advisor Central Home page that is displayed when you click the Advisor Central link under the Related Links heading on the Database Home page, or on other pages where it is listed. You can also invoke advisors in the context of recommendations from alerts. Advisors are powerful tools for tuning your database. In general, advisors produce more comprehensive recommendations than alerts, since alert generation is intended to be low cost and have minimal impact on performance. On the other hand, because Advisors are often user-invoked, they can consume more resources and perform more detailed analysis. Their automated analysis can provide more results than you might normally be able to produce manually in the time afforded you as part of regular operations. This analysis, along with the what-if capability of some advisors, provides vital information for tuning that cannot be procured from any other source.

## Masking Sensitive Data for Non-production Use

Data masking (also known as data scrambling and data anonymization,) is the process of replacing sensitive information copied from production databases to test or non-production databases with realistic, but scrubbed, data based on masking rules. Data masking is ideal for virtually any situation when confidential or regulated data needs to be shared with other non-production users; for instance, internal users such as application developers, or external business partners, like offshore testing companies or suppliers and customers. These non-production users need to access some of the original data, but do not need to see every column of every table, especially when the information is protected by government regulations.

Data masking allows organizations to generate realistic and fully functional data with similar characteristics as the original data to replace sensitive or confidential information. This contrasts with encryption or Virtual Private Database, which simply hides data, and the original data can be retrieved with the appropriate access or key. With data masking, the original sensitive data cannot be retrieved or accessed.

Names, addresses, phone numbers, and credit card details are examples of data that require protection of the information content from inappropriate visibility. Live production database environments contain valuable and confidential data — access to this information is tightly controlled. However, each production system usually has replicated development copies, and the controls on such test environments are less stringent. This greatly increases the risks that the data might be used inappropriately. Data masking can modify sensitive database records so that they remain usable, but contain no confidential or personally identifiable information. Yet, the masked test data resembles the original in appearance to ensure the integrity of the application.

## Security and Regulatory Compliance

Masked data is a sensible precaution from a business security standpoint, because masked test information can help prevent accidental data escapes. In many cases, masked data is a legal obligation. The Enterprise Manager Data Masking Pack can help organizations fulfill legal obligations and comply with global regulatory requirements, such as Sarbanes-Oxley, the California Database Security Breach Notification Act (CA Senate Bill 1386), and the European Union Data Protection Directive.

The legal requirements vary from country to country, but most countries now have regulations of some form to protect the confidentiality and integrity of personal consumer information. For example, in the United States, The Right to Financial Privacy Act of 1978 creates statutory Fourth Amendment protection for financial records, and a host of individual state laws require this. Similarly, the U.S. Health Insurance Portability and Accountability Act (HIPAA) created protection of personal medical information.

## Roles of Data Masking Users

The following types of users participate in the data masking process for a typical enterprise:

- Application database administrator or application developer  
This user is knowledgeable about the application and database objects. This user may add additional custom database objects or extensions to packaged applications, such as the Oracle E-Business suite.
- Information security administrator  
This user defines information security policies, enforces security best practices, and also recommends the data to be hidden and protected.

## Format Libraries and Masking Definitions

To mask data, the Data Masking Pack provides two main features:

- **Masking format libraries**  
The format library contains a collection of ready-to-use masking formats. The library consists of format routines that you can use for masking. A masking format can either be one that you create, or one from the list of Oracle-supplied default masking formats.  
  
As a matter of best practice, organizations should create masking formats for all common regulated information so that they can be applied to the sensitive data regardless of which database the sensitive data resides in. This ensures that all sensitive data is consistently masked across the entire organization.
- **Masking definitions**  
A masking definition defines a data masking operation to be implemented on one or more tables in a database. Masking definitions associate table columns with formats to use for masking the data. They also maintain the relationship between columns that are not formally declared in the database using related columns.

To create a masking definition, you specify the column of the table for which the data should be masked and the format of masked data. If the columns being masked are involved in unique, primary key, or foreign key constraints, data masking generates the values so that the constraints are not violated. You can create a new masking definition or use an existing definition for a masking operation. You would typically export masking definitions to files and import them on other systems. This is important when the test and production sites reside on different Oracle Management Systems or reside on entirely different sites.

Data masking is an iterative and evolving process handled by the security administrator and orchestrated by the database administrator. When you first configure data masking, try out the masking definition on a test system, then add a greater number of columns to the masking definition and test it to make sure it functions correctly and does not break any application constraints. During this



process, you should exercise care when removing all imbedded references to the real data while maintaining referential integrity. After data masking is configured to your satisfaction, you can use the existing definition to repeatedly mask after cloning. The masking definition, however, would need to evolve as new schema changes require new data and columns to be masked.

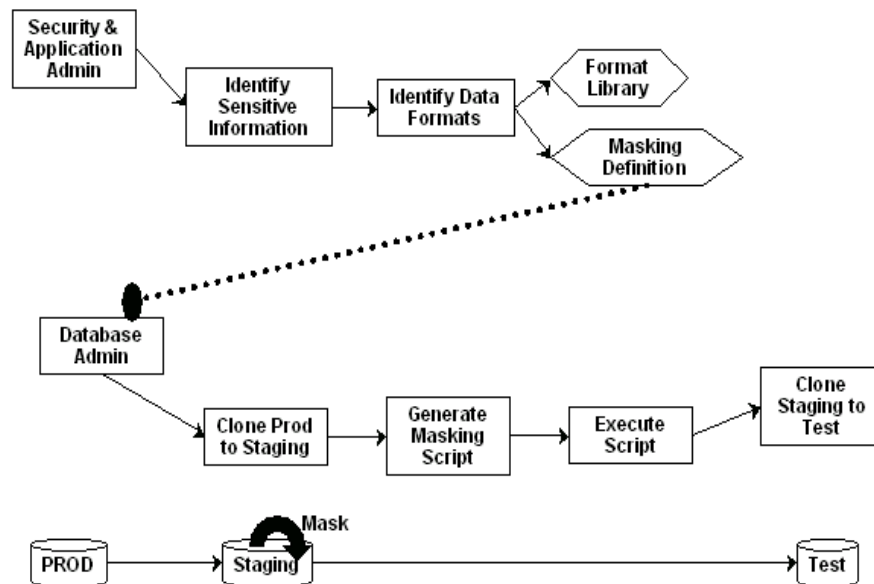
**See Also:** "Creating a Data Masking Definition" in the Enterprise Manager online help as well as the help for each Data Masking page.

## Recommended Data Masking Workflow

Figure 5–13 shows that the production database is cloned to a staging region and then masked there. During the masking process, the staging and test areas are tightly controlled like a production site.

After the masking process is complete, you can distribute the database for wide availability. If you need to ship the database to another third-party site, you are required to use the Data Pump Export utility, and then ship the dump file to the remote site. However, if you are retaining the masked data in-house, see the next section.

**Figure 5–13 Data Masking Workflow**



## Data Masking Task Sequence

The task sequence in this section demonstrates the data masking workflow and refers you to additional information elsewhere in this chapter for some of the tasks in the sequence. Before reviewing this sequence, note that there are two options for completing this process:

- You can clone the production database to a staging area, mask it, then export/import it to another database before delivering it to in-house testers or external customers. This is the most secure approach.

- You can clone the production database to a mask staging area, then make the staging area the new test region. In this case, you should not grant testers SYSDBA access or access to the database files. Doing so would compromise security. The masked database contains the original data in unused blocks and in the free list. You can only purge this information by exporting/importing the data to another database.

The following steps guide you through the data masking process. Note the references to other sections for supporting information.

1. Review the application database and identify the sources of sensitive information.
2. Define mask formats for the sensitive data. The mask formats may be simple or complex depending on the information security needs of the organization.

See "[Data Masking Task Sequence](#)" on page 5-23 for more information.

3. Create a masking definition to associate table columns to these mask formats. Data masking determines the database foreign key relationships and adds foreign key columns to the mask.

See "[Creating a Masking Definition](#)" on page 5-27 for more information.

4. Optionally declare dependent columns not defined in the database but enforced by the applications. Masking assumes it should perform the masking to honor these additional constraints.

This requires knowledge of the application schema. Consult the application documentation to identify the relationship between the tables and the columns containing sensitive data to ensure complete coverage of your application data.

See "[Adding Dependent Columns](#)" on page 5-30 for more information.

5. Save the masking definition and generate the masking script.
6. Clone the production database to a staging area, selecting the masking definition to be used after cloning. Note that you can perform cloning using Enterprise Manager, which enables you to add masking to the Enterprise Manager clone workflow. However, if you performing cloning outside of Enterprise Manager, you must initiate masking from Enterprise Manager after cloning is complete. The cloned database should be controlled with the same privileges as the production system since it still contains sensitive production data.

After cloning, be sure to change the passwords as well as update or disable any database links, streams, or references to external data sources. Back up the cloned database or, at a minimum, the tables that contain masked data. This will help you restore the original data if the masking definition needs to be refined further.

See "[Cloning the Production Database](#)" on page 5-30 for more information.

7. Verify if the masked data meets the information security requirements. Otherwise, refine the masking definition, restore the altered tables, and reapply the masking definition until the optimal set of masking definitions has been identified.
8. After masking, test all of your applications, reports, and business processes to ensure they are functional. If everything is working, you can export the masking definition to keep it as a backup.
9. After masking the staging site, make sure to drop any tables named MGMT\_DM\_ TT before cloning to a test region. These temporary tables contain a mapping between the original sensitive column value and the mask values and are therefore sensitive in nature.

During masking, Enterprise Manager automatically drops these temporary tables for you with the default "Drop temporary tables created during masking" option. However, you can preserve these temporary tables by unselecting this option. In this case, you are responsible for deleting the temporary tables before cloning to the test region.

10. After masking is complete, ensure that all tables loaded for use by the substitute column format or table column format will be dropped. These tables contain the mask values that table column or substitute formats will use. It is recommended that you purge this information for security reasons.  
See "[Deterministic Masking](#)" on page 5-30 for more information.
11. Clone the database to a test region or use it as the new test region. When cloning the database to an external or unsecured site, you should use Export or Import, and only supply the data in the database rather than the database files themselves.
12. Provide the masking definition to the application database administrator to use in masking the database as part of cloning production for testing.

## Defining New Masking Formats

A masking definition requires one or more masking formats for any columns included in the masking definition. When adding columns to a masking definition, you can either create masking formats manually or import them from the format library. It is often more efficient to work with masking formats from the format library.

To create a masking format in the format library, do the following:

1. Navigate to the **Databases** targets roll-up page. The Databases page appears.
2. Click the **Data Masking Format Library** link. The Format Library appears with predefined formats that Oracle Enterprise Manager provides.
3. Click **Create**. The Create Format page appears, where you can define a masking format.

**Tip:** For information on page user controls, see the online help for the Format page.

4. Provide a required name for the new format, select a format entry type from the **Add** drop-down list, then click **Go**.

A page appears that enables you to provide input for the format entry you have selected. For instance, if you select Array List, the subsequent page enables you to enter a list of values, such as New York, New Jersey, and New Hampshire.

5. Continue adding additional format entries as needed.
6. When done, provide an optional user-defined or post-processing function (see "[Providing User-defined and Post-processing Functions](#)"), then click **OK** to save the masking format.

The Format Library page reappears with your newly created format displayed in the Format Library table.

**Tip:** For information on page user controls, see the online help for the Format Library and Create Format pages.

### Providing User-defined and Post-processing Functions

If desired, you can provide user-defined and post-processing functions on the Create Format page. A user-defined choice is available in the Add drop-down, and post-processing function field is available at the bottom of the page.

#### ■ User-defined functions

A user-defined function passes in the original value as input and returns a mask value. The data type and uniqueness of the output values must be compatible with the original output values. Otherwise, a failure occurs when the job runs. Combinable, a user-defined function is a PL/SQL function that can be invoked in a Select statement. Its signature is returned as:

```
Function udf_func (rowid varchar2, column_name varchar2, original_value  
varchar2) returns varchar2;
```

- rowid is the min (rowid) of the rows that contain the value original\_value 3rd argument.
- column\_name is the name of the column being masked.
- original\_value is the value being masked.

That is, it accepts an input string that is the original value and returns the mask value.

Both input and output values are varchar2. For instance, a user-defined function to mask a number could receive 100 as input, the string representation of the number 100, and return 99, the string representation of the number 99. Values are cast appropriately when inserting to the table. If the value is not castable, masking fails.

#### ■ Post-processing functions

A post-processing function has the same signature as a user-defined function, but passes in the mask value the masking engine generates, and returns the mask value that should be used for masking, as shown in the following example:

```
Function post_proc_udf_func (rowid varchar2, column_name varchar2, mask_value  
varchar2) returns varchar2;
```

- rowid is the min (rowid) of the rows that contain the value mask\_value.
- column\_name is the name of the column being masked.
- mask\_value is the value being masked.

### Using Masking Format Templates

After you have created at least one format, you can use the format definition as a template in the Create Format page, where you can implement most of the format under a different name and change the entries as needed, rather than having to create a new format from scratch.

To create a new format that is similar to an existing format, select a format on the Format Library page and click **Create Like**. The masking format you select can either be one you have previously defined yourself, or one from the list of out-of-box masking formats. You can use these generic masking format definitions for different applications.

See "[Using Oracle-supplied Predefined Masking Formats](#)" for instructional details about the various Oracle-supplied predefined masking format definitions and how to modify them to suit your needs.

## Using Oracle-supplied Predefined Masking Formats

The following sections discuss the various Oracle-supplied format definitions and how to modify them to suit your needs.

### Patterns of Format Definitions

All of the format definitions adhere to these typical patterns:

- Generate a random number or random digits.
- Perform post-processing on the above-generated value to ensure that the final result is a valid, realistic value

For example, a valid credit card number must pass Luhn's check. That is, the last digit of any credit card number is a checksum digit, which is always computed. Also, the first few digits indicate the card type (MasterCard, Amex, Visa, and so forth). Consequently, the format definition of a credit card would be as follows:

- Generate random and unique 10-digit numbers.
- Using a post-processing function, transform the values above to a proper credit card number by adding well known card type prefixes and computing the last digit.

This format is capable of generating 10 billion unique credit card numbers.

### Category Definitions

The following sections discuss different categories of these definitions. By default, these mask formats are also available in different format styles, such as a hyphen ( - ) format. If needed, you can modify the format style.

#### Credit Card Numbers

Out of the box, the format library provides many different formats for credit cards. The credit card numbers generated by these format pass the standard credit card validation tests by the applications, thereby making them appear like valid credit card numbers.

Some of the credit card formats you can use are:

- MasterCard numbers
- Visa card numbers
- American Express card numbers
- Discover Card numbers
- Any credit card number  
(Credit card numbers belong to all types of cards.)

You may want to use different styles for storing credit card numbers, such as:

- Pure numbers
- 'Space' for every four digits
- 'Hyphen' ( - ) for every four digits, and so forth

To implement the masked values in a certain format style, you can set the `DM_CC_FORMAT` variable of the `DM_FMTLIB` package. To install the `DM_FMTLIB` package, see ["Installing the DM\\_FMTLIB Package"](#) on page 5-26.

### **United States Social Security Numbers**

Out of the box, you can generate valid U.S. social security (SSN) numbers. These SSNs pass the normal application tests of a valid SSN.

You can affect the format style by setting the `DM_SSN_FORMAT` variable of the `DM_FMTLIB` package. For example, if you set this variable to '-', the typical social security number would appear as '123-45-6789'.

### **ISBN Numbers**

Using the format library, you can generate either 10-digit or 13-digit ISBN numbers. These numbers adhere to standard ISBN number validation tests. All of these ISBN numbers are random in nature. Similar to other format definitions, you can affect the "style" of the ISBN format by setting values to `DM_ISBN_FORMAT`.

### **UPC Numbers**

Using the format library, you can generate valid UPC numbers. They adhere to standard tests for valid UPC numbers. You can affect the formatting style by setting the `DM_UPC_FORMAT` value of the `DM_FMTLIB` package.

### **Canadian Social Insurance Numbers**

Using the format library, you can generate valid Canadian Social Insurance Numbers (SINs). These numbers adhere to standard tests of Canadian SINs. You can affect the formatting style by setting the `DM_CN_SIN_FORMAT` value of the `DM_FMTLIB` package.

### **North American Phone Numbers**

Out of the box, the format library provides various possible U.S. and Canadian phone numbers. These are valid, realistic looking numbers than can pass standard phone number validation tests employed by applications. You can generate the following types of numbers:

- Any North American phone numbers
- Any Canadian phone number
- Any U.S.A. phone number

### **Installing the DM\_FMTLIB Package**

The predefined masking formats use functions defined in the `DM_FMTLIB` package. This package is automatically installed in the `DBSNMP` schema of your Enterprise Manager repository database. To use the predefined masking formats on a target database (other than the repository database), you must manually install the `DM_FMTLIB` package on that database. To do this, follow these instructions:

1. Locate the following scripts in your Enterprise Manager installation:

```
$ORACLE_HOME/sysman/admin/emdrep/sql/db/latest/masking/dm_fmtlib_pkgdef.sql  
$ORACLE_HOME/sysman/admin/emdrep/sql/db/latest/masking/dm_fmtlib_pkgbody.plb
```

2. Copy these scripts to a directory in your target database installation and execute them using `SQL*Plus`, connected as a user that can create packages in the `DBSNMP` schema.

After you have done this, you can use the predefined masking formats in your masking definition.

## Creating a Masking Definition

Before creating a masking definition, be aware of the following advisory information:

- Make sure the format you select does not violate check constraints and does not break any applications that use the data.
- For triggers and PL/SQL packages, data masking recompiles the object.
- Exercise caution when masking partitioned tables, especially if you are masking the partition key. In this circumstance, the row may move to another partition.
- Data Masking does not support clustered tables, masking information in object tables, XML tables, and virtual columns. Relational tables are supported for the masking.
- If objects are layered on top of a table such as views, materialized views, and PL/SQL packages, they are recompiled to be valid.

To create a masking definition, do the following:

1. Navigate to the **Databases** targets roll-up page. The Databases page appears.
2. Click the **Data Masking Definitions** link under the Related Links section. The Data Masking Definitions page appears, where you can create and schedule new masking definitions and manage existing masking definitions.

**Tip:** For information on page user controls, see the online help for the Data Masking Definitions page.

3. Click **Create** to go to the Create Masking Definition page. A masking definition includes information regarding table columns and the format for each column. You can choose which columns to mask, leaving the remaining columns intact.

**Tip:** For information on page user controls, see the online help for the Create Masking Definition page.

4. Specify a required database name. The Database field shows the database name of the database target you selected before accessing the Data Masking Definitions Page. You can specify any of the available databases. To search for the list, click the flashlight icon.
5. Click **Add** to go to the Add Columns page, where you can add one or more columns for masking and automatically add foreign key columns. You need to add at least one column in the masking definition.

**Tip:** For information on page user controls, see the online help for the Add Columns page.

6. In the Search section, minimally provide a schema name and optionally provide input for the other search fields to filter the results, then click **Search**.
7. Either select one or more columns for formatting later on the Create Masking Definition page, or formatting now if the data types of the columns you have selected are identical.
8. Decide if you want to mask selected columns as a group.

Select this checkbox if you want to mask more than one column together, rather than separately. When you then select two or more columns and then later define the format on the Define Group Mask page, the columns appear together, and any

choices you make for format type or masking table apply collectively to all of the columns.

Columns that you want to mask as a group must all be from the same table.

9. Either click **Add** to add the column to the masking definition, return to the Create Masking Definition page, and define the format of the column later, or click **Define Format and Add** to define the format for the column now.

The Define Format and Add feature can save you significant time. When you select multiple columns to add that have the same data type, you do not need to define the format for each column as you would when you click Add. For instance, if you search for social security numbers (SSN) and this yields 100 SSN columns, you could select them all, then click Define Format and Add to import the SSN format for all of them.

10. **If you clicked Add in the previous step...**

... you will eventually need to define the format of the column in the Create Masking Definition page before you can proceed. When you are ready to do so, click the icon in the page **Format** column for the column you want to format. Depending on whether you decided to mask selected columns as a group on the Add Columns page, either the Define Column mask or Define Group mask appears. Read further in this step for instructions for both cases.

**If you clicked Define Format and Add in the previous step and did not check "Mask selected columns as a group," ...**

... the Define Column Mask page appears, where you can define the format for the column before adding the column to the Create Masking Definition page.

- a. Provide a format entry for the required Default condition by either selecting a format entry from the drop down list and clicking **Add**, or clicking **Import Format**, selecting a predefined format on the Import Format page, then clicking **Import**.

See "[Using Oracle-supplied Predefined Masking Formats](#)" for information about Oracle-supplied predefined masking format definitions.

**Tip:** For information on page user controls, see the online help for the Define Column Mask page.

- b. Add another condition by clicking **Add Condition** to add a new condition row, then provide one or more format entries as described in the previous step. [Figure 5-14](#) shows an example of user input.



Figure 5–14 Sample Input for Define Column Mask Page

ORACLE Enterprise Manager 10g  
Grid Control

Home Targets Deployments Alerts Compliance Jobs Reports

Databases Hosts Services Systems Groups All Targets

Define Column Mask

Owner HR Column NATIONAL\_ID  
Table EMPLOYEES Data Type VARCHAR2(20)

Cancel OK

By default all records in the table will be masked using the specified format. You can optionally identify more than one subset of records using conditions. Each subset can be masked using a corresponding masking format. The subsets will be masked in the order they are specified. A subset will not be masked again even when it matches a subsequent condition.

Import Format Format Entry Post Processing Function Add

Expand All Collapse All

Select	Condition	Format Entry Properties				Sample	Remove
		Property	Value	Property	Value		
<input type="checkbox"/>	▼ COUNTRY='US'					123-45-6789	
	Random Digits	Start Length	3	End Length	3		
	Fixed String	String	-				
	Random Digits	Start Length	2	End Length	2		
	Fixed String	String	-				
	Random Digits	Start Length	4	End Length	4		
<input type="checkbox"/>	▶ COUNTRY='CA'					123-456-789	
<input type="checkbox"/>	▼ COUNTRY='EN'						
	Preserve Original Data						
	Post Processing Function	Function Name	HR.MSK				
<input type="checkbox"/>	▼ Default						
	Drop Data Rows						

- c. When you have finished formatting the column, click **OK** to return to the Create Masking Definition page.

If you clicked **Define Format and Add** in the previous step and checked "Mask selected columns as a group," ....

... the Define Group Mask page appears, where you can add format entries for group columns that appear in the Create Masking Definition page.

- a. Select one of the available format types. For complete information on the format types, see the online help.

**Tip:** For information on page user controls, see the online help for the Define Group Mask page.

- b. Optionally add a column to the group.
- c. When you have finished formatting the group, click **OK** to return to the Create Masking Definition page. Your configuration appears in the Columns table.

11. Click **OK** to save your definition and return to the Data Masking Definitions page.

12. Select the definition and click **Generate Script** to view the script for the list of database commands used to mask the columns you selected earlier.

From the Script Generation Results page that appears, you can clone and mask the database or schedule the data masking job.

**Tip:** For information on page user controls, see the online help for the Script Generation Results page.

13. Make one of the following choices:

- If you are working with a production database, click **Clone and Mask**.
- If you are already working with a test database and want to directly mask the data in this database, click **Schedule Job**.

## Adding Dependent Columns

The following prerequisites apply for the column to be defined as dependent:

- A valid dependent column should not already be included for masking.
- The column should not be a foreign key column or referenced by a foreign key column.
- The column data should conform to that in the parent column.

If the column does not meet these criteria, an "Invalid Dependent Columns" message appears when you attempt to add the dependent column.

To add dependent columns, do the following:

1. From the Create Masking Definition page, click the **+ icon** in the Dependent Columns Add column for the column you want to format. The Add Dependent Columns page appears.

**Tip:** For information on page user controls, see the online help for the Add Dependent Columns page.

2. In the Search section, minimally provide a schema name or select it by first clicking on the flashlight search icon, then clicking **Search**.
3. Select one or more column names from the list that appears below, then click **Add**.  
The Create Masking Definition page reappears and displays your added dependent column in the Dependent Columns table at the bottom of the page. The dependent column will be masked using the same format as that specified for the parent column.

## Cloning the Production Database

To clone and optionally mask the masking definition's target database, do the following:

1. Ensure that you have a Provisioning and Patch Automation pack license before proceeding. The Clone Database feature requires this license.
2. From the Data Masking Definitions page, select the masking definition you want to clone, then click **Clone Database**. The Clone Database wizard appears, where you can create a test system to run the mask.
3. Proceed through the wizard steps as you ordinarily would to clone a database. For assistance, refer to the online help for each step.
4. In the Database Configuration step of the wizard, add a masking definition.
5. Schedule and then run the clone job.

## Deterministic Masking

You may occasionally need to consistently mask multiple, distinct databases. For instance, if you run HR, payroll, and benefits that have an employee ID concept on three separate databases, the concept may be consistent for all of these databases, in

that an employee's ID can be selected to retrieve the employee's HR, payroll, or benefits information. Based on this premise, if you were to mask the employee's ID because it actually contains his/her social security number, you would have to mask this consistently across these three databases.

Deterministic masking provides a solution for this problem. You can use the Substitute format to mask employee ID column(s) in all three databases. The Substitute format uses a table of values from which to substitute the original value with a mask value. As long as this table of values does not change, the mask is deterministic or consistent across the three databases.

**Tip:** For more information on the Substitute format, see the online help for the Define Column Mask page.

## Use Cases

The following use cases show some of the more important data masking tasks that you can perform.

### **Mask production data for an application maintained in a staging environment**

This use case applies to the application database administrator except for step 6.

1. Log into the system as the application database administrator to access databases and their associated privileges.
2. Search the schema for tables that contain columns with confidential or regulated information. The system responds by displaying tables and columns based on search criteria.
3. Select the table columns as well as the appropriate mask formats. The system responds by adding the table column and all of its associated foreign key columns into the masking definition.
4. Assign additional table columns to the masking definition wherein explicit referential integrity relationships are not defined in the database. The system responds by adding these table columns as dependent columns to the masking definition.
5. Notify the information security administrator of the masking definition.
6. The information security administrator reviews the masking definition, verifies its completeness, then recommends that the application database administrator proceed with the masking process.
7. Initiate the data masking process using the masking definition created in step 4. The system responds by validating the masking definition, verifying space availability, then applying the mask formats to the data in the test-staging database.
8. Verify the completeness of the masking process, and confirm that masked data has replaced sensitive data.

### **Use an existing masking definition to mask a new database**

This use case applies to the application database administrator.

1. Create a new masking definition based on a previously created masking definition that is then applied to a new database. The system responds by copying the masking definition, including table columns and mask formats, and associating it with the new database.

2. Use the newly created masking definition to initiate masking of the new database. The system responds by validating the masking definition, verifying space availability, then applying the mask formats to the data in the test-staging database.

### **Export and import masking definitions**

This use case applies to the application developer and application database administrator.

1. The developer logs into the system to access databases and their associated privileges.
2. The developer requests the export of a valid masking definition into a transportable file format. The system responds by generating an XML file with the masking definition and mask formats.
3. The database administrator downloads or receives mask templates.
4. The database administrator imports the mask template. The system responds by validating the mask template and creating a masking definition based on the template.
5. The database administrator uses the newly created masking definition to initiate masking of the database. The system responds by validating the masking definition, verifying space availability, then applying the mask formats to the data in the test-staging database.

## **Related Oracle Security Offerings**

Besides data masking, Oracle offers the following security products:

- **Virtual Private Database or Oracle Label Security** — Hides rows and data depending on user access grants.
- **Transparent Data Encryption** — Hides information stored on disk using encryption. Clients see unencrypted information.
- **DBMS\_CRYPTO** — Provides server packages that enable you to encrypt user data.
- **Database Vault** — Provides greater access controls on data.

## **Maintaining Databases**

You can use Oracle Enterprise Manager to control the flow of data between or outside Oracle Databases. The following sections provide an overview of the functions available to you that can help you maintain your Oracle Database.

### **Using Backup**

Backup of an Oracle Database generally refers to physical backup; protecting the files that make up your database. The files protected by the backup and recovery facilities built into Oracle Enterprise Manager include datafiles, control files, and archived redo log files. The backup mechanisms that work at the physical level protect against damage at the file level, such as the accidental deletion of a datafile or the failure of a disk drive. The focus in Oracle backup and recovery is generally on the physical backup of database files, which permit the full reconstruction of your database.

Enterprise Manager's physical backup and recovery features are built on Oracle's Recovery Manager (RMAN) command-line client. Enterprise Manager carries out its

backup tasks by composing RMAN commands and sending them to the RMAN client. Enterprise Manager makes available much of the functionality of RMAN, as well as wizards and automatic strategies to simplify and further automate implementing RMAN-based backup and recovery.

Managing backups consists of two things: managing the backups themselves as they exist on disk or tape, and managing the record of backups kept in the RMAN repository.

Datafiles or archived redo logs copied at the operating system level can be cataloged and added to the RMAN repository, after which they can be used in data restore and recovery operations just as if they had been created with RMAN. Backup maintenance provided in Enterprise Manager includes the following:

- Viewing lists of backups (backup sets and image copies) recorded in the RMAN repository
- Cross-checking your repository, which marks as expired any backups that are in the repository but not accessible at the time of the crosscheck
- Deleting expired backups from your RMAN repository
- Deleting obsolete backups from the repository and from disk. Note that if you use a recovery area for your backup storage, many maintenance activities are reduced or eliminated because of the flash recovery area's automatic management of disk space.

## Recovery

Media recovery using Enterprise Manager can be either complete recovery or point-in-time recovery. In complete recovery, all changes from the logs are applied and the database returns to its state at the time of failure. You can then reopen the database with no loss of data.

In point-in-time recovery, you can choose any System Change Number (SCN) between the time of the datafile backup and the last change in the redo logs, and apply only changes up through that SCN. You can thus return your database to any SCN (and thus any point-in-time) between the time of your backup and the most recent SCN in the redo logs. This technique can be used to recover from situations such as user errors that cause logical corruption in the database. Point-in-time recovery is sometimes referred to as incomplete recovery, since not all changes are applied.

Media recovery requires a control file, datafiles, and all online and archived redo logs from the time the datafiles were backed up. It is typically used only in the case of database failure.

Crash recovery is used to recover from a failure either when a single-instance database crashes, or all instances of an Oracle RAC database crash. Instance recovery refers to the case where a surviving instance recovers a failed instance in an Oracle RAC database. Datafile media recovery is used to recover from a lost or damaged current datafile or control file. Block media recovery is a technique for restoring and recovering individual data blocks while all database files remain online and available.

## Flashback Recovery

Enterprise Manager's flashback features provide a range of physical and logical data recovery tools as efficient, easy-to-use alternatives to physical and logical backups. Flashback table allows you to revert a table to its contents at a time in the recent past; and flashback drop allows you to rescue dropped database tables. Neither requires

advance preparation such as creating logical-level exports to allow for retrieval of your lost data, and both can be used while your database is available.

Flashback database lets you quickly recover an Oracle Database to a previous time to correct problems caused by logical data corruptions or user errors. If a flash recovery area is configured, then you can return the database to a prior time. Flashback table lets you recover tables to a specified point in time. You can restore table data along with all its associated attributes, such as indexes, triggers, and so on. This is done while the database is online by rolling back only the changes to the given tables. You can revert the table and its contents to a specific time or user-specified SCN. Use flashback table with flashback query and row versions to find a time to which the table should be restored.

The larger the flash recovery area, the more useful it becomes. Ideally, the flash recovery area should be large enough to hold two complete backup copies of your datafiles, plus any incremental backups and archive logs required to restore your database to any point in time during your recovery window.

## Backup and Recovery Settings

You can use Enterprise Manager to configure a number of settings and policies that determine how backups are stored, which data is backed up, how backups perform, and how long backups are retained before being purged from the recovery area. You can also configure settings to improve backup performance. For disks, you can configure the default format for storing backups, the location on disk where backups are stored, and whether backup tasks are run in parallel for improved performance. Oracle backups to disk can be saved as image copies or backup sets. Backups to tape and similar media management devices must be stored as backup sets.

## Managing Restore Points

You can use Enterprise Manager to create a restore point as a designated point in time to which you can restore your database. A restore point is a name associated with a past point-in-time of the database. You can flash back the database to a restore point if the required flashback logs and archived logs exist. A Guaranteed Restore Point is a restore point to which you can always flash back the database. Each restore point has a name and creation time. The restore points are sorted by creation time beginning with the most recent.

## Overview of Data Guard

Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data. Data Guard provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable production Oracle Databases to survive disasters and data corruptions. Data Guard maintains these standby databases as transactionally consistent copies of the production database. Then, if the production database becomes unavailable because of a planned or an unplanned outage, Data Guard can switch any standby database to the production role, thus minimizing the downtime associated with the outage.

Data Guard can be used with traditional backup, restoration, and cluster techniques to provide a high level of data protection and data availability. With Data Guard, administrators can optionally improve production database performance by offloading resource-intensive backup and reporting operations to standby systems.

A Data Guard configuration consists of one production database and one or more standby databases. You can manage primary and standby databases using Enterprise

Manager. You can use Enterprise Manager to monitor the status of a configuration as well as the online redo log file activity of the primary and standby databases. At the most basic level, the Data Guard Overview page for the configuration not only displays information about the configuration, but it also includes summary information about its databases.

## Additional Maintenance Features

Enterprise Manager allows you to easily move data from files or databases into an existing database. You can export and import data, and you can clone databases using the tools available in Enterprise Manager. The following sections describe the features available in the Maintenance area of the Database Home page.

### Exporting and Importing Features

The Export to Files feature enables you to move existing data in Oracle format to and from Oracle Databases. For example, export files can archive database data or move data among different Oracle Databases that run on the same or different operating systems. Using Export to Files, you can back up logical database objects while the database is open and available for use. It writes a read-consistent view of the database's objects to an operating system file.

Conversely, you can use Enterprise Manager to import the contents of a database, objects, and tables. You can also import the contents of a database by using the Import from Database feature. You can import an entire database, schemas within a database and the objects in the schemas, or one or more tables from within a schema. Use the Load Data from File feature to load data from a non-Oracle database into an Oracle Database.

Use the Monitor Export and Import Job feature to view the status of an import and export operation such as a full database export or a tablespace export. You can change the state of a job by suspending it, canceling it, or resuming it if it had previously been suspended. You can also increase the number of threads dedicated to the job and thus increase its resources.

### Cloning Databases or Transporting Tablespaces

You can use the Enterprise Manager Clone Database tool to clone an Oracle Database instance to an existing Oracle home. After you have an Oracle Database instance in a known state, you may want to clone that database to another existing Oracle home.

Use the Transport Tablespaces feature to transport tablespaces between different machine architectures and operating systems. Transportable tablespaces entirely bypass the unload and reload steps. Using transportable tablespaces, Oracle datafiles (containing table data, indexes, and almost every other Oracle Database object) can be directly transported from one database to another. You can use the transportable tablespaces feature to move a subset of an Oracle Database and plug it in to another Oracle Database, essentially moving tablespaces between the databases.

### Overview of Oracle Streams

Oracle Streams enables the propagation and management of data, transactions, and events in a data stream either within a database, or from one database to another. The stream routes published information to subscribed destinations. As users' needs change, they can simply implement a new capability of Oracle Streams, without sacrificing existing capabilities. Oracle Streams can stream data between databases,

nodes, or blade farms in a grid and can keep two or more copies in sync as updates are applied. It also provides a unified framework for information sharing, combining message queuing, replication, events, data warehouse loading, notifications, and publish/subscribe into a single technology.

Oracle Streams provides a set of elements that lets users control what information is put into a stream, how the stream flows or is routed from node to node, what happens to events in the stream as they flow into each node, and how the stream terminates. By specifying the configuration of the elements acting on the stream, you can address specific requirements, such as message queuing or data replication.

Use the tools in Oracle Enterprise Manager to set up and manage your Oracle Streams environment.

## Working with Software Configurations

The Oracle Management Agent on a host collects host configuration information for the host and database configuration information for the Oracle Databases on the host and client configuration information and communicates that information over HTTPS to the Oracle Management Service, which stores it in the Oracle Management Repository. Enterprise Manager lets you compare these configurations to determine differences between two or more hosts, clients or databases. The Generic Compare feature allows you to compare various types of current/saved configurations with one or more current/saved configurations.

Using this feature, you can compare the current configuration of a selected target type with one or more current configurations of other targets of the same type, or compare saved configurations with one or more saved configurations of the same or other targets. You can also compare saved configurations with one or more current configurations of the same or other targets, or compare a specific configuration with another configuration and list the differences immediately. Lastly, you can compare a specific configuration with another configuration and schedule the comparison as a job.

## Using Database Software Patching

Use Enterprise Manager to simplify the patching of Oracle software on any host where an Oracle Management Agent is running, and to provide critical patch advisories. Enterprise Manager simplifies the process of patching Oracle software. Oracle Patch Advisories describe critical software patches for Oracle products. To help ensure a secure and reliable configuration, all relevant and current critical Oracle patches should be applied. To promote critical patch application, Enterprise Manager performs an assessment of vulnerabilities by examining the host configurations collected for your enterprise to determine the Oracle homes that require one or more critical patches to be installed. All the Critical Patch Advisories are listed with their corresponding Impact areas, a brief description of each advisory, the number of Affected Hosts, and Oracle homes for each advisory.

You can connect to *My Oracle Support* via Enterprise Manager, perform a search, download the required patches/patch sets, and apply the patches. You can perform all the patching activities from the Patch Cache. That is, even when the OMS is not connected to *My Oracle Support* via the Internet, you can perform a search, download, and apply a patch/patch set.



## Monitoring Oracle Real Application Clusters

Oracle Real Application Clusters (RAC) provides a high-availability database environment spanning multiple hosts. Each cluster can be comprised of multiple cluster databases, each of which consists of multiple cluster database instances. A cluster database is available as long as one of its instances is available.

Enterprise Manager provides performance pages to monitor all levels of a cluster environment, including the cluster, the cluster database, and the cluster database instances. Managing Oracle Real Application Clusters databases and instances is similar to managing single-instance databases. Using Enterprise Manager, you can perform various tasks, such as:

- Managing clusters, cluster databases, and cluster database instances
- Monitoring key performance statistics
- Performing administration and maintenance tasks
- Managing Oracle Real Application Clusters features, such as Cluster Managed Database Services

Oracle RAC enables each computer (or host) that is a member of the cluster to share access to the database. If one cluster host fails or is taken offline, the other hosts of the cluster continue operating, and the entire Oracle RAC database remains available for applications. This means that two or more computers with typical performance appear to applications as if they were a much more powerful computer.

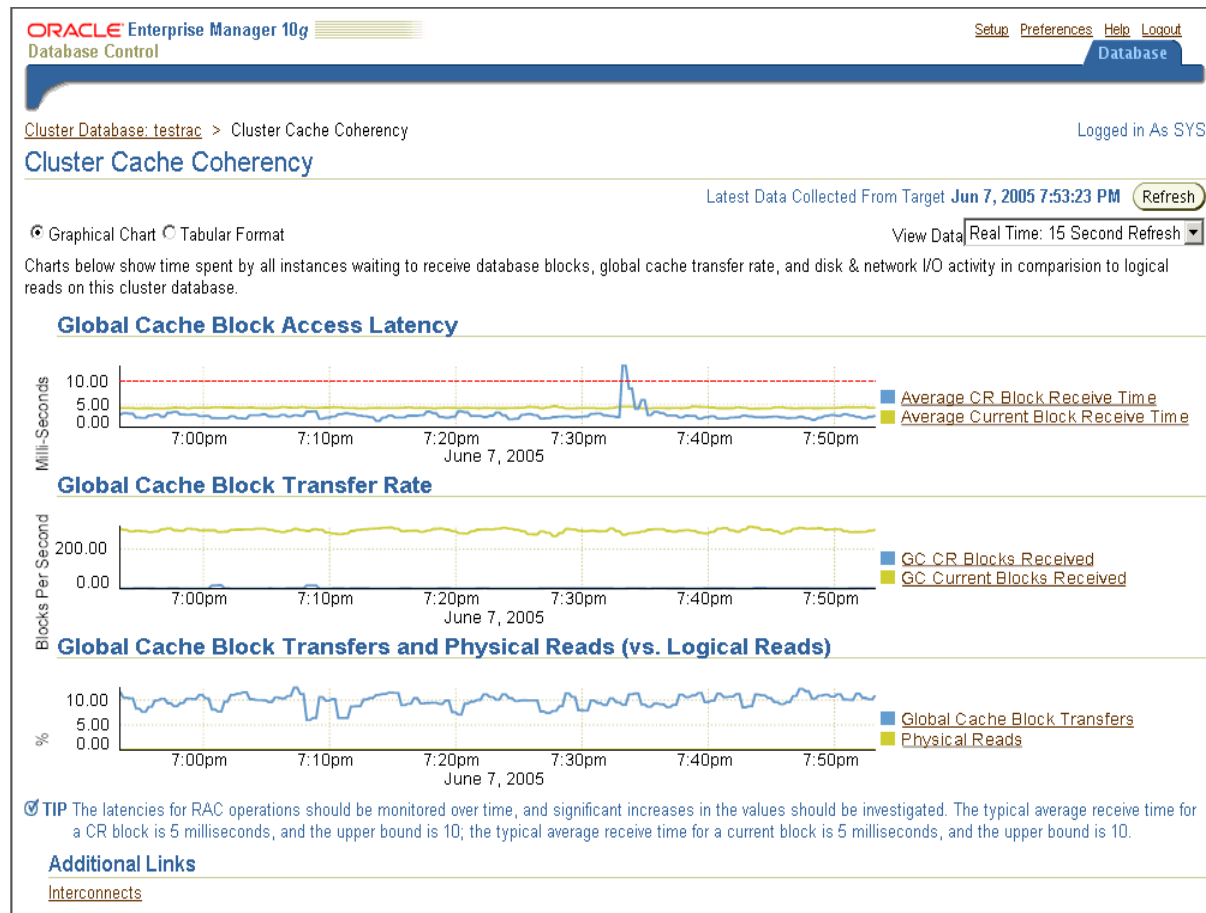
To increase performance, availability, and reliability of a two-host Oracle RAC database, you can add cluster hosts. Because data is not partitioned between hosts, adding hosts to the cluster does not create instability; instead, applications can run faster or support more users. The more hosts your Oracle RAC database has, the less the loss of any individual node affects the database.

### Cluster Cache Coherency

Concurrent read and write activity on shared data in a cluster occurs frequently. Depending on the service requirements, this activity does not normally cause performance problems. However, when global cache requests cause a performance problem as indicated on the Cluster Database page, a successful strategy for performance tuning is to optimize SQL plans and the schema to achieve effective local cache hit efficiency and minimize I/O. To assist you in resolving the problem, the Cluster Cache Coherency page enables you to view cache coherency metrics for the entire cluster database, identify processing trends, and optimize performance for your Oracle Real Application Clusters environment (Figure 5-15).

**See Also:** "Cluster Cache Coherency Page" in the Enterprise Manager online help

Figure 5–15 Cluster Cache Coherency Page



The charts have the following purposes:

- **Global Cache Block Access Latency** represents the end-to-end elapsed time or latency for a block request.
- **Global Cache Block Transfer Rate** shows the total aggregated number of data blocks received by all instances in the cluster through an interconnect.
- **Global Cache Block Transfers and Physical Reads** shows the percentage of logical reads that read data from the buffer cache of other instances through Direct Memory Access and from disk.

## Cluster Interconnects

The Cluster Interconnects page enables you to view the current state of interfaces on hosts (Figure 5–16). You can use this page to monitor the interconnect interfaces, determine configuration issues, and identify transfer rate-related issues, such as excess traffic. This page helps determine the load added by individual instances and databases on the interconnect. Sometimes, you can immediately identify interconnect delays due to applications outside the Oracle Database.

**See Also:** "Cluster Interconnects Page" in the Enterprise Manager online help

Figure 5–16 Cluster Interconnects Page

ORACLE Enterprise Manager 10g Database Control Setup Preferences Help Logout Database

Cluster: crs Latest Data Collected From Target Jun 7, 2005 7:48:09 PM Refresh

Home Performance Targets **Interconnects**

The interconnect configuration and internode communication will influence the performance of cluster databases. The tables below show network interfaces on all hosts and network interfaces currently in use by cluster databases. It is important that cluster databases are configured to use a private interconnect for message and block transfers.

Private Interconnect Transfer Rate (MB/Sec) **3.21**  
Transfer rate on the private network in the last 5 minutes.

### Interfaces by Hosts

View

Name	Type	Subnet	Interface Type	Total I/O Rate (MB/Sec) (Last 5 Minutes)	Total Error Rate (%) (Last 5 Minutes)
crs	Cluster				
staju11.us.oracle.com	Host				
eth2	Interface	n/a	Unknown	n/a	n/a
eth0	Interface	10.87.24.0	Private	3.18	0
eth1	Interface	140.87.24.0	Public	1.54	0
staju12.us.oracle.com	Host				
eth0	Interface	10.87.24.0	Private	3.23	0
eth2	Interface	n/a	Unknown	n/a	n/a
eth1	Interface	140.87.24.0	Public	1.27	0

### Interfaces in Use by Cluster Databases

Name	Target Type	Interface Name	Host Name	IP Address	Interface Type	Source	Transfer Rate (MB/Sec) (Last 5 Minutes)
testrac	Cluster Database						
testrac1	Database Instance	eth0	staju11.us.oracle.com	10.87.25.147	Private	Oracle Cluster Repository	2.5
testrac2	Database Instance	eth0	staju12.us.oracle.com	10.87.25.148	Private	Oracle Cluster Repository	2.5

TIP The Transfer Rate is the estimated traffic contributed by the instance assuming uniform block size in the database.

Home Performance Targets **Interconnects**

### Hosts

Name	Status	Alerts	Policy Violations	CPU Util %	Mem Util %	Total IO/sec
staju11.us.oracle.com		1 0	23	19.14	99.61	17.91
staju12.us.oracle.com		0 0	18	13.24	43.19	11.35

## Cluster-Managed Database Services

Services are groups or classifications of applications that comprise business components that extend and respond to application workloads. Examples of services are Accounts Payable, Customer Relationship Management, and so on. Services in Oracle RAC enable continuous, uninterrupted database operations to support multiple services on multiple instances.

Services enable Oracle RAC to integrate cluster database resources into a single system image to optimize cluster manageability. This simplifies system deployment, testing, disaster recovery, and administrative overhead. With services, users connect to a database without regard for which instance executes the SQL application service.

You assign services to run on one or more instances, and alternate instances can serve as backup instances in case the primary instance fails. If a primary instance fails, Oracle moves the service from the failed instance to a surviving alternate instance.

Services enable you to model and deploy both planned and unplanned operations for all types of high availability or disaster recovery scenarios. During outages, Oracle RAC automatically restarts key components. Components that are eligible for automatic restart include instances, Oracle Net Services listeners, and the database as well as several database subcomponents.

You can create and edit services using the Create Services page in Enterprise Manager, which you can access from the Cluster Managed Database Services link in the Enterprise Manager Maintenance page (Figure 5–17).

**See Also:** "Create Service Page" in the Enterprise Manager online help

**Figure 5–17 Create Service Page**

## Oracle Clusterware and High Availability

When you combine Oracle Clusterware and Oracle RAC, you can achieve excellent scalability and high availability. To maintain high availability, the Oracle Clusterware components can respond to status changes to restart applications and processes according to high availability rules you can specify in the Create Service page (Figure 5–17). Oracle Clusterware achieves high availability with the components in Table 5–2.

**Table 5–2 Oracle Clusterware High Availability Components**

Component	Description
Voting Disk	The voting disk manages cluster membership by using a health check, and arbitrates cluster ownership among the instances in case of network failures. Oracle RAC uses the voting disk to determine which instances are members of a cluster.

**Table 5–2 (Cont.) Oracle Clusterware High Availability Components**

<b>Component</b>	<b>Description</b>
Oracle Clusterware Registry	The OCR maintains database and cluster database configuration information as well as configuration information about any cluster database within the cluster. The OCR also manages information about processes that Oracle Clusterware controls.
Application Programming Interface	Oracle Clusterware provides a high availability application programming interface (API) that you can use to manage applications or processes that run on single-instance Oracle databases or Oracle RAC databases. This enables you to provide high availability for all of your applications.

## Database High Availability

The Database High Availability Console is a one-stop, dashboard style page for monitoring your database high availability (HA). You can use the High Availability Console page to display high availability events, view the availability summary that includes the status of the database, view the last backup status, view the Flash Recovery Area Usage, view the Data Guard summary, and view the RAC Services summary including Top Services. You can view the Database High Availability Console page in Advanced View or Basic View, depending on how much detail you want to display.

You can use the Database High Availability Console Options page to set up event options for availability events, configure formatting options for the header bar and target and host names, define refresh options on the High Availability Console page, and specify whether and how to display historical charts.

You can use the Database Maximum Availability Architecture (MAA) Advisor page to view recommended Oracle solutions for each outage type (computer failures, storage failures, human errors, data corruptions, and site failures), view configuration status and Enterprise Manager links for each solution, and view the benefits for each of the solutions that are offered. The MAA Advisor page contains a table that lists the Outage Type, Oracle Solution, Recommendation Level, Configuration Status, and Benefits. You can navigate to a page where you can set up, manage and configure the specific solution area. Once a solution has been configured, you can click on the Advisor Details link to see the updated values.



---

---

## Service Management

This chapter introduces the concepts of services and service management in the following sections.

- [Introduction to Services](#)
- [Modeling Services](#)
- [Managing Systems](#)
- [Monitoring Services](#)
- [Diagnosing Service Problems](#)

### Introduction to Services

The critical and complex nature of today's business applications has made it very important for IT organizations to monitor and manage application service levels at high standards of availability. Problems faced in an enterprise include service failures and performance degradation. Since these services form an important type of business delivery, monitoring these services and quickly correcting problems before they can impact business operations is crucial in any enterprise.

Service-level agreements are used to evaluate service availability, performance, and usage. By constantly monitoring the service levels, IT organizations can identify problems and their potential impact, diagnose root causes of service failure, and fix these in compliance with the service-level agreements.

Enterprise Manager provides a comprehensive monitoring solution that helps you to effectively manage services from the overview level to the individual component level. When a service fails or performs poorly, Grid Control provides diagnostics tools that help to resolve problems quickly and efficiently, significantly reducing administrative costs spent on problem identification and resolution. Finally, customized reports offer a valuable mechanism to analyze the behavior of the applications over time.

Grid Control monitors not only individual components in the IT infrastructure, but also the applications hosted by those components, allowing you to model and monitor business functions using a top-down approach, or from an end-user perspective. If modeled correctly, services can provide an accurate measure of the availability, performance, and usage of the function or application they are modeling.

### Defining Services in Enterprise Manager

A "service" is defined as an entity that provides a useful function to its users. Some examples of services include CRM applications, online banking, and e-mail services.

Some simpler forms of services are business functions that are supported by protocols such as DNS, LDAP, POP, or SMTP.

Grid Control allows you to define one or more services that represent the business functions or applications that run in your enterprise. You can define these services by creating one or more service tests that simulate common end-user functionality. Using these service tests, you can measure the performance and availability of critical business functions, receive alerts when there is a problem, identify common issues, and diagnose causes of failures.

You can define the following service types: Generic Service, Web Application, Forms Application, and Aggregate Service. Web applications, a special type of service, are used to monitor Web transactions. Forms applications are used to model and monitor Forms transactions.

The following elements are important to understanding Grid Control's Service Level Management feature:

- **Service:** Models a business process or application.
- **Availability:** A condition that determines whether the service is considered accessible by the users or not.
- **Service Test:** The functional test defined by the Enterprise Manager administrator against the service to determine whether or not the service is available and performing.
- **System:** A group of underlying components, such as hosts, databases, and application servers, on which the service runs. For more information on systems, refer to the "[Managing Systems](#)" section in this chapter.
- **Beacons:** A functionality built into Management Agents used to pre-record transactions or service tests.
- **Performance and Usage:** Performance indicates the response time as experienced by the end users. Usage refers to the user demand or load on the system.
- **Service Level:** Operational or contractual objective for service availability and performance.
- **Root Cause Analysis:** Diagnostic tool to help determine the possible cause of service failure

## Modeling Services

You can create a new target, called a service, to model and monitor your business applications from within Grid Control. While creating a service, you can define the availability, performance and usage parameters, and service-level rules.

### Availability

"Availability" of a service is a measure of the end users' ability to access the service at a given point in time. However, the rules of what constitutes availability may differ from one application to another. For example, for a Customer Relationship Management (CRM) application, availability may mean that a user can successfully log on to the application and access a sales report. For an online store, availability may be monitored based on whether the user can successfully log in, browse the store, and make an online purchase.

Grid Control allows you to define the availability of your service based on service tests or systems.



- **Service Test-Based Availability:** Choose this option if the availability of your service is determined by the availability of a critical functionality to your end users. Examples of critical functions include accessing e-mail, generating a sales report, performing online banking transactions, and so on. While defining a service test, choose the protocol that most closely matches the critical functionality of your business process, and beacon locations that match the locations of your user communities. You can define one or more service tests using standard protocols and designate one or more service tests as "Key Tests." These key tests can be executed by one or more "Key Beacons" in different user communities. A service is considered available if one or all key tests can be executed successfully by at least one beacon, depending on your availability definition.
- **System-Based Availability:** Your service's availability can alternatively be based on the underlying system that hosts the service. Select the components that are critical to running your service and designate one or more components as "Key Components," which are used to determine the availability of the service. The service is considered available as long as at least one or all key components are up and running, depending on your availability definition. For more information on systems, refer to the "[Managing Systems](#)" section in this chapter.

## Performance and Usage

You can define metrics to measure the performance and usage of the service. Performance indicates the response time of the service as experienced by the end user. Usage metrics are based on the user demand or load on the system.

Performance metrics are collected for service tests when the service tests are run by beacons. You can calculate the minimum, maximum, and average response data collected by two or more beacons. For example, you can monitor the time required to retrieve e-mails from your e-mail service in San Francisco, Tokyo, and London, then compare results. You can also collect performance metrics for system components, then calculate the minimum, maximum, and average values across all components. For example, you can monitor average CPU utilization, memory utilization, and disk I/O utilization across several hosts.

Usage metrics are collected based on the usage of the system components on which the service is hosted. For example, if you are defining an e-mail service that depends on an IMAP server, you can use the Total Client Connections metric of the IMAP server to represent the usage of this e-mail service. You can monitor the usage of a specific component or statistically calculate the minimum, maximum, and average values from a set of components. You can also set thresholds on the above metrics and receive notifications and alerts. For more information on setting thresholds, see the "[Monitoring Services](#)" section of this chapter.

## Business Metrics

You can define metrics to measure the performance of business in an organization. Business metrics are collected for system based services. Non system based metrics can be defined by using the Data Exchange feature. For more details on the Data Exchange feature, refer to *Oracle Enterprise Manager Integration Guide*.

.You can define the business metric to be based on the performance of one specific system component or aggregate it across multiple components. You can also set thresholds on these metrics and receive notifications and alerts.

## Setting Service-Level Rules

Service-level parameters are used to measure the quality of the service. These parameters are usually based on actual service-level agreements or on operational objectives.

Grid Control's Service Level Management feature allows you to proactively monitor your enterprise against your service-level agreements to verify that you are meeting the availability, performance, and business needs within the service's business hours. For service-level agreements, you may want to specify the levels according to operational or contractual objectives.

By monitoring against service levels, you can ensure the quality and compliance of your business processes and applications.

## Monitoring Templates for Services

Administrators are often faced with the task of defining similar monitoring attributes or rules for many applications. The same set of rules are often applicable to different applications. This can be achieved through the Monitoring Templates feature in Grid Control. A monitoring template for a service contains definitions for one or more service tests, as well as a list of monitoring beacons. You can create a monitoring template from a standard service target, then copy this template to create service tests for any number of service targets and specify a list of monitoring beacons. This helps reduce the required configuration time where a large number of applications need to be monitored.

## Managing Systems

A **"system"** is a logical grouping of targets that collectively hosts one or more services. It is a set of infrastructure targets (hosts, databases, application servers, and so on) that function together to host one or more applications or services.

In Enterprise Manager, systems constitute a new target type. For example, to monitor an e-mail application in Enterprise Manager, you would first create a system, such as "Mail System," that consists of the database, listener, application server, and host targets on which the e-mail application runs. You would then create a service target to represent the e-mail application and specify that it runs on the Mail System target.

---

---

**Note:** An Enterprise Manager "System" is used specifically to monitor the components on which a service runs. Many of the functions and capabilities for groups and systems are similar.

---

---

## Creating Systems

Use the Create System pages to perform the following configuration tasks:

- Select target components for a new system.
- Define the associations between the components of the system using the Topology Viewer.
- Add charts that will appear in the System Charts page. The charts represent the overall performance for the system or components of the system. Based on the target type of the components you select in the Components page, some charts are predefined.

- Select a set of columns you want to appear in the System Components page and in the system's Oracle Grid Control Dashboard.
- Customize the refresh frequency and specify the format for viewing component status, alerts, and policy violations in the system's Oracle Grid Control Dashboard.

Enterprise Manager provides a Topology Viewer for several applications. The Topology Viewer allows you to view the relationships between components, nodes, or objects within different Oracle applications. You can zoom, pan, see selection details and summary information, and evaluate aggregate components. Individually distinct icons are used for each object type, and standardized visual indicators are used across all applications.

You may want to create system topologies for a number of reasons:

- Graphically model relationships
- Identify the source of a failure
- Perform visual analysis for high-level problem detection

When creating a system topology, you specify associations between the components in the system to logically represent the connections or interactions between them. For example, you can define an association between the database and the listener to indicate the relationship between them. Components are represented as icons, and associations are depicted as arrow links between components. After you have customized the topology to suit your needs, you can then view the overall status of the components in your system by accessing the System Topology page.

**See Also:** "Create System Topology Page" in the Enterprise Manager online help


## Monitoring Systems

Use the System pages to perform the following monitoring and administration tasks:

- Quickly view key information about components of a system, such as outstanding alerts and policy violations.
- View metric data for several time periods.
- View summary information determined by the columns you configured when you created the system.
- Perform administrative tasks, such as creating jobs and blackouts.
- View the topology of system components, including the associations between them.

**See Also:** "[System Topology](#)" in this chapter

Figure 6-1 System Home Page



## Monitoring Services

Monitoring a service helps you ensure that your operational and service-level goals are met. To monitor a service, define service tests that simulate activity or functionality that is commonly accessed by end users of the service. For example, you may want to measure a service based on a particular protocol, such as DNS, LDAP, and IMAP. To proactively monitor the availability and responsiveness of your service from different user locations, designate the geographical locations from which these service tests will be executed. Run service tests from specified locations using Enterprise Manager Beacons. You may also measure a service based on the usage of the service's system components.

## Services Dashboard

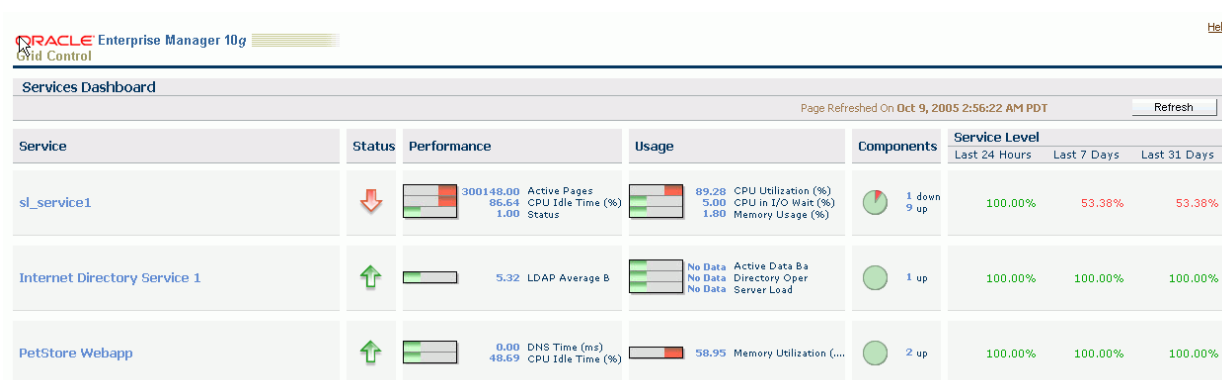
In Grid Control, service levels are defined as the percentage of time during business hours that a service meets the specified availability, performance, and business criteria. Using the Services Dashboard, administrators can determine whether the service levels are compliant with business expectations and goals.

The Services Dashboard enables administrators to browse through all service-level-related information from a central location. The Services Dashboard illustrates the availability status of each service, performance and usage data, as well as service-level statistics. You can easily drill down to the root cause of the problem or determine the impact of a failed component on the service itself.

The following details are displayed in the Services Dashboard:

- **Availability:** A measure of the end users' ability to access the service at a given point in time. Service level agreements typically require a service be available at least for a minimum percentage of time.
- **Performance:** Response time is a good measure of the performance experienced by the end users when they access the service. When the service performance is poor, the availability of the service may be affected.
- **Usage:** Indicates end-user usage, or level of user activity, of the service.
- **Business:** Business metrics measure the performance of business in an organization. Business metrics are displayed only if one or more key business indicators are associated with the service. The Service Level is impacted when a critical alert is generated for the specified business metric.


**Figure 6–2 Consolidated Dashboard View of Multiple Services**



## System Topology

The System Topology page enables you to view the dependency relationships between components of the system. From the topology view, you can drill down to detail pages to get more information on the key components, alerts and policy violations, possible root causes and services impacted, and more.

Use the System Topology page, shown in [Figure 6–3](#), to get a quick overview of the status of your system's components. The status indicators over each icon enable you to quickly assess which components are down or have open alerts. You can get more detailed information for any key component from this page.

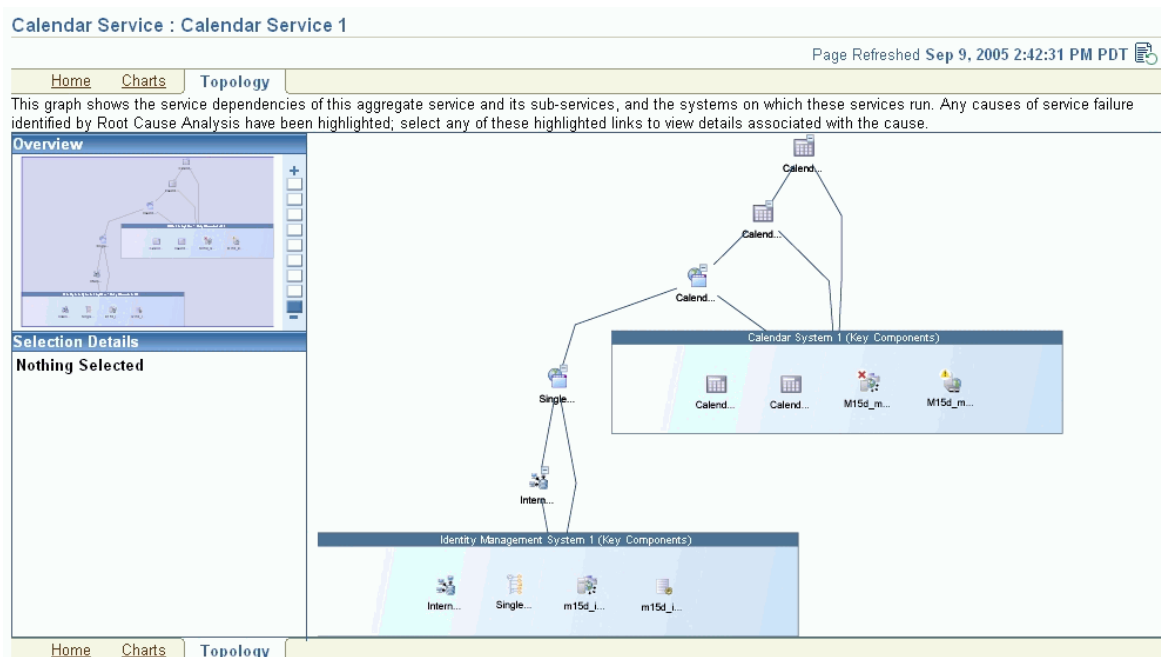
**Figure 6–3 System Topology Page**

**See Also:** ["Creating Systems"](#) in this chapter

## Service Topology

Use the Service Topology page, shown in [Figure 6–4](#), to view the dependencies between the service, its system components, and other services that define its availability. Upon service failure, the potential causes of failure, as identified by Root Cause Analysis, are highlighted in the topology view. In the topology, you can view dependent relationships between services and systems.

Some data centers have systems dedicated to one application or service, while others have shared systems that host multiple services. In Grid Control, you can associate a single service or multiple services with a system, based on the setup of the data center.

**Figure 6–4 Service Topology Page for a Collaboration Suite Service**

## Reports

Enterprise Manager provides out-of-box reports that are useful for monitoring services and Web applications. You can also set the publishing options for reports so that they are sent out via email at a specified period of time. Some of the reports that can be generated include Web Application Alerts, Web Application Transaction Performance Details, and Service Status Summary.

## Notifications, Alerts, and Baselines

Using Grid Control, you can proactively monitor a service and address problems before users are impacted. Each service definition has performance and usage metrics that have corresponding critical and warning thresholds. When a threshold is reached, Grid Control displays an alert. There are a standard set of notification rules that specify the alert conditions for which notifications should be sent to the appropriate administrators. Apart from these standard sets of rules, you can define and set up schedules so that administrators are notified when the specified alerts conditions are met. For example, thresholds can be defined so that alerts are generated when a system is down, if the end user cannot login to an application, or if the online transaction cannot be successfully completed.

You can set up baselines for a specified period and use these baselines to evaluate performance. Statistics are computed over the baseline period for specific target metrics. You can use these statistics to automatically set metric thresholds for alerting, as well as to normalize graphical displays of service performance.

## Service Performance

Grid Control provides a graphical representation of the historic and current performance and usage trends in the Performance and Usage Charts. You can view metric data for the current day (24 hours), 7 days, or 31 days. The thresholds for any performance or usage alerts generated during the selected period are also displayed in

the charts. This helps you to easily track the performance and usage of the service test or system over time and investigate causes of service failure. Users can choose the default chart for the Services Home page; all performance and usage charts are available on the Charts page.

Use the Test Performance page to view the historical and current performance of the service tests from each of the beacons. If a service test has been defined for this service, then the response time measurements as a result of executing that service test can be used as a basis for the service's performance metrics. It is possible to have multiple response time measurements if the service access involves multiple steps or the service provides multiple business functions. Alternatively, performance metrics from the underlying system components can also be used to measure performance of a service.

If performance of a service seems slow, it may be due to high usage of the service. Monitoring the service usage helps diagnose poor performance by indicating whether the service is affected by high usage of a system component.

## Monitoring Web Application Services

Today's e-businesses depend heavily upon their Web applications to allow critical business processes to be performed online. As more emphasis is placed on accessing information quickly, remotely, and accurately, how can you ensure your online customers can successfully complete a transaction? Are you certain that your sales force is able to access the information they need to be effective in the field?

The Web application management features complement the traditional target monitoring capabilities of Enterprise Manager. Full integration with the Enterprise Manager target monitoring capabilities allow you to monitor the performance and availability of components that make up the applications' technology environment, including the back-end database and the middle-tier application servers.

In Grid Control, you can define a Web application service to monitor Web transactions. This allows you to proactively monitor your e-business systems from the top down, and trace the experience of your end users as they enter and navigate the Web site. You can monitor the Web application service through the Services Dashboard, Topology Viewer, Charts, Reports, and more. For more information on these features, refer to the "[Monitoring Services](#)" section in this chapter.

Additionally, you can monitor the end-user performance response times, which enables you to effectively manage your e-business systems and understand the impact of application service-level problems.

### Transactions

Transactions are service tests that are used to test the Web application performance and availability. Important business activities for the Web application are recorded as transactions, which are used to test availability and performance of a Web application. A transaction is considered "available" if it can be successfully executed by at least one beacon. You can record the transaction using an intuitive playback recorder that automatically records a series of user actions and navigation paths.

### End-User Performance Monitoring

The End-User Performance Monitoring feature enables you to measure the actual response time as experienced by the end users. When configured with Oracle Application Server Web Cache or Oracle HTTP Server / Apache HTTP Server, the End-User Performance Monitoring feature provides response time data generated by actual end users as they access and navigate your Web site.



You can track the response times for each user and all individual pages, allowing you to assess the end-user experience and address potential issues. You can also view the response times by individual visitor, domain, user-defined region, Web server, or a combination of these criteria. For example, tracking the response time of visitors ensures that critical customers, executives, and other important visitors are experiencing adequate response times.

You can set up Watch Lists of important URLs and view the response metrics of these critical pages at a glance. You can also use the Analyze feature to analyze the performance data stored in the Management Repository.

## Monitoring Forms Applications

Oracle Forms provides a set of integrated builders that allow application developers to easily and quickly construct sophisticated database forms and business logic with minimal effort. The powerful, declarative features in Oracle Forms enable developers to create fully functional applications from database definitions with significantly reduced coding.

A Form is a collection of UI widgets that serves a specific purpose. It is usually built to perform one isolated task (such as accessing and displaying user data). A Forms application is a set of Forms that are related to each other. For example, Financials is a Forms application that contains Forms like, `DISPLAY_SALARY`, `EDIT_ACCOUNT`, and `PRINT_TAX_ACCOUNT`.

A Forms Application target in Enterprise Manager is used to model and monitor a specific Forms application. You can define, monitor, and playback a specific flow in the Forms application where each flow is recorded as a single Forms transaction. You can also measure the response time data as experienced by end-users while they are performing a Forms operation. You can measure the total response time, server time, and the database time for Forms operations such as Commit, Query, Runform, Callform, Newform, and Openform

### Forms Transactions

A Forms transaction consists of a set of user actions within a single application when using Forms. You can record multiple Forms transactions by using the intuitive playback recorder that automatically records the Forms actions. You can monitor these transactions periodically and collect relevant metrics when these transactions are played back.

### End-User Performance Monitoring

The End-User Performance Monitoring utility allows you to measure the response time of your applications by viewing information about how quickly the responses are. When configured with Oracle Application Server Web Cache or Oracle HTTP Server/Apache HTTP Server, the End-User Performance Monitoring feature provides response time data generated by actual end users as they access a Forms application and perform a set of Forms actions.

When you access a Forms application, the End-User Performance Monitoring utility measures the response time of Forms actions such as Commit, Query, Runform, Callform, Newform, and Openform. You can monitor the Forms actions and view reports based on the response times experienced by the user. You can also define a Watch List of the most important Forms actions to monitor and view the response metrics of these critical operations at a glance.

## Diagnosing Service Problems

Grid Control offers you tools to help diagnose service problems, including Root Cause Analysis, Topology Viewer, and Web application diagnostics. If a service is unavailable or performing poorly, use these tools to determine the potential causes.

### Root Cause Analysis

When a service fails, Root Cause Analysis returns a list of potential causes on the Service Home page. Potential root causes include failed subservices and failed key system components.

By default, Root Cause Analysis evaluates a key component's availability status to determine whether or not it is a cause of service failure. You can specify additional conditions, or component tests, for Root Cause Analysis to consider. If a key component is unavailable, or if any of your component test's conditions are not met, then this component is considered a possible cause of the service failure.

You can also specify additional conditions, or component host tests, for the host on which this key component resides. If Root Cause Analysis identifies the key component as a cause of service failure, the component's host is then analyzed to see if it potentially caused the component, and therefore the service, to fail.

You can also access the Root Cause Analysis information from the Topology Viewer, which shows a graphical representation of the hierarchical levels displaying relationships between components. Red lines between the services and system components represent the associated failure. Follow these red lines to discover possible causes of failure.

Grid Control can also be integrated with the EMC SMARTS solution to detect network failures in Root Cause Analysis. When problems in the network are detected, you can use the SMARTS network adapter to query Root Cause Analysis information related to the hosts and IP addresses in the network.

### Diagnosing Web Application Problems

When a Web application is unavailable, the Root Cause Analysis feature allows you to determine the causes of service failure. Apart from this feature, Grid Control provides tools to diagnose application performance degradation issues and pinpoint problem areas within the application stack. Comprehensive diagnostic tools enable you quickly drill down into the Oracle Application Server stack and monitor response times in various application server and database components.

#### Interactive Transaction Tracing

When the performance of a Web application is slow, you can trace problematic transactions as required using Interactive Transaction Tracing. You can record the transaction using an intuitive playback recorder that automatically records a series of user actions and navigation paths. You can play back transactions interactively and perform an in-depth analysis of the response times across all tiers of the Web application for quick diagnosis.

The Interactive Transaction Tracing facility complements the Transaction Performance Monitoring and End-User Performance Monitoring features by helping you diagnose the cause of a performance problem. This in-depth drill-down diagnostics tool enables you to trace the transaction path and performance across the application tiers, and helps identify the cause of performance bottlenecks. Using these diagnostic tools, you can quickly resolve application problems, thus reducing the mean-time to repair.

All invocation paths of a transaction are traced and hierarchically broken down by servlet/JSP, EJB, and database times to help you locate and solve the problem quickly. Once a problem is resolved, you can also run Interactive Transaction Tracing to reassure you that the problem has been satisfactorily repaired. In addition, you can use the **SQL Statement Analysis** link to view details.

### **Request Performance Diagnostics**

Grid Control provides in-depth historical details on the J2EE and database performance of all URL requests. By examining the detailed J2EE and database breakdown and analyzing the processing time of a request, you can determine whether the problem lies within a servlet, JSP, EJB method, or specific SQL statement. Using this information, you can easily isolate the cause of the problem and take necessary action to quickly repair the appropriate components of your Web application.

Grid Control's Request Performance Diagnostics feature is instrumental to the application server and back-end problem diagnosis process. Slowest URL request processing times and the number of hits are provided so that you can easily recognize where problem resolution efforts should be prioritized. Application administrators need to know how their J2EE and database components are performing, including the top JSPs and servlets by processing time and request rates so that they can identify how these components are affecting overall response times.

URL request processing time and load graphs provide you with information on the impact of server activity on response times. Analyzing the J2EE and database at the subcomponent level helps you make accurate decisions to tune or repair the appropriate elements of a Web application.

Easy to read graphs of URL request processing times by the OC4J subsystem allows you to quickly assess where the most time is spent. Further drill-downs bring you directly to in-depth URL request processing call stack details. You can correlate URL request times (EJB time, database time, and so on) to the underlying system component metrics.



---

---

# Change Management for Databases

This chapter introduces the concept of change management in the following sections:

- [Overview of Change Management for Databases](#)
- [Dictionary Baselines](#)
- [Dictionary Comparisons](#)
- [Dictionary Synchronizations](#)

## Overview of Change Management for Databases

To manage the lifecycle of enterprise applications, an organization will need to maintain multiple copies of an application database for various purposes such as development, staging, production, and testing. Each of these databases must adhere to different processes. For example, for production databases, it is essential to ensure adherence to proper production control procedures. It is vital that administrators have the tools to detect unauthorized changes, such as an index being dropped without the requisite change approvals. In such cases, monitoring changes to production databases day over day or week over week becomes vital.

Database compliance, that is, ensuring that all databases meet the gold standard configuration, is another important aspect of life cycle management. Compliance with organizational standards or best practices ensures database efficiency, maintenance, and ease of operation.

On development databases, developers make changes that the database administrator needs to consolidate and propagate to staging or test databases. The goal is to identify the changes made to development and then make the same changes to staging or test databases taking into account any other changes already in production database.

Typically, most applications will get upgraded over time. Also, most applications are customized by the business user to suit their needs. Application customizations are usually dependent on database objects or PL/SQL modules supplied by the application vendor. The application vendor supplies the upgrade scripts and the customer has very little transparency about the impact of the upgrade procedure on their customizations. When customers test upgrade databases, they can capture a baseline of the application schema before and after the upgrade. A comparison of the before and after baselines will tell the user what modules were changed by the application. This gives them a better idea about how their customizations will be impacted as a result of upgrading their application.

There are three core capabilities of Change Management that allow developers and database administrators to manage changes in database environments:

- **Baseline** - A point in time of the definition of the database and its associated database objects.
- **Comparison** - A complete list of differences between a baseline or a database and another baseline or a database.
- **Synchronization** - The process of promoting changes from a database definition capture in a baseline or from a database to a target database.

For database versions 9.x and above, the user logged into the database target through Enterprise Manager must have `SELECT ANY DICTIONARY` privilege for capturing or comparing databases. In addition to the `SELECT ANY DICTIONARY` privilege, the user logging into the destination database for creating dictionary synchronization needs to be a database administrator (DBA) or must have appropriate privileges on the objects being synchronized.

## Dictionary Baselines

A dictionary baseline contains a set of database definitions captured at a certain point in time. Baselines are stored in the Enterprise Manager repository, and are in the form of an XML document called Simple XML or SXML.

Each baseline must be assigned a unique name. A good practice to name baselines is to match it on the scope of the database objects being captured in the baseline, for example, Financial 11.5.10 or HR Benefits or PO Check Print. A baseline can have a series of versions that have been captured at different points in time. Creating multiple versions of a baseline allows you to track changes to the definitions of a set of database objects over time. You can compare two versions of the same baseline to see the changes that have occurred between them.

When creating a baseline, you also create a corresponding baseline scope specification, which describes the names and the types of database objects and schemas from which they should be captured. When you have created the baseline definition, you can then capture the first version of the baseline by submitting an Enterprise Manager job. At a later time, or at regular intervals, you can capture additional versions of the same baseline. Each baseline version records the metadata definitions as they exist at the time the version is captured.

## Scope Specification

A scope specification identifies the database objects to be captured in a baseline. (Scope specifications also identify objects to process in dictionary comparisons and synchronizations.) Once you have specified the scope of a baseline, you cannot change the scope specification. This restriction ensures that all versions of the baseline are captured using the same set of rules, which means that differences between versions result from changes in the database, not scope specification changes. To capture dictionary objects using a different scope specification, you must create a new baseline.

Baseline scope specifications take three forms.

- You can specify schemas and object types to capture. For example, you can capture all Tables, Indexes and Views in schemas APPL1 and APPL2. This form of scope specification is appropriate when there is a well-defined set of schemas that contain your application objects. In addition to schema objects, you can also capture non-schema objects (such as Users, Roles and Tablespaces) and privilege grants to Users and Roles.

- You can specify schemas to exclude, and object types. This form of scope specification captures objects that are contained in all schemas other than those you specify. For example, you can capture all object types in schemas other than SYSTEM and SYS. This form of scope specification is appropriate when you want to capture all database objects, with the exception of objects contained in Oracle-provided schemas. As with the first form of scope specification, you can also capture non-schema objects and privilege grants.
- Finally, you can capture individual schema objects by specifying the type, schema and name of each object. This form of scope specification is appropriate when you want to capture a few specific objects, rather than all objects contained within one or more schemas. While capturing individual schema objects, you can also capture non-schema objects and privilege grants.

If you include a non-schema object type, such as User or Role, in a scope specification, all objects of that type are captured. There is no way to capture individual non-schema objects. [Figure 7-1](#) shows the Dictionary Baselines:Objects page where the scope can be specified.

**Figure 7-1** Baseline Scope Specification in the form Of Schemas And Object Types to Capture

ORACLE Enterprise Manager 10g  
Grid Control

Home | Targets | Deployments | Alerts | Compliance | Jobs | Reports

Hosts | Databases | Web Applications | Services | Systems | Groups | All Targets

Source | **Objects** | Job Options | Review

### Create Dictionary Baseline: Objects

Source Database: appdb1  
Logged In As: system

Specify which database objects you want to capture.

#### Non-Schema Object Types

Tablespace     Rollback Segment     User     Role     Profile

Grants for Users and Roles

#### Schema Objects

Object Selection Method: Select object types and the schemas to include

Object Types

Select All | Select None

Table     Index     View     Trigger     Sequence

Package     Package Body     Procedure     Function     Cluster

Database Link     Materialized View     Materialized View Log     Private Synonym     Public Synonym

User Defined Type

#### Schemas to Include

Remove All | Add

Schema	Remove
HR	
SCOTT	

▶ Object Name Prefix

#### Database Attributes

Initialization Parameters

Cancel | Back | Step 2 of 4 | Next

Home | Targets | Deployments | Alerts | Compliance | Jobs | Reports | Setup | Preferences | Help | Logout

Copyright © 1996, 2009, Oracle and/or its affiliates. All rights reserved.  
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.  
Other names may be trademarks of their respective owners.  
[About Oracle Enterprise Manager](#)

## Capturing a Dictionary Baseline Version

As the final step of defining a baseline, you specify when to capture the first version of the baseline. You can capture the first version immediately, or at a later time (for example, when the database is not being used in active development work). You can also indicate that additional versions of the baseline should be captured at regular intervals without further intervention on your part.

You can also capture a new baseline version at any time by selecting the baseline and specifying "Recapture Now."

Baselines processed after the initial version generally complete substantially faster than the initial version. Only those objects that have changed are captured in the new version. This also means that the storage required for additional baseline versions is only slightly larger than the storage used by the initial version, assuming a small percentage of objects has changed. [Figure 7-2](#) shows the first version of a dictionary baseline.

**Figure 7-2 A Baseline Viewed After Completion of the Capture of the First Version**

ORACLE Enterprise Manager 10g

Grid Control

Home Targets Deployments Alerts Compliance Jobs Reports

Hosts Databases Web Applications Services Systems Groups All Targets

Dictionary Baselines >

Baseline: my baseline Refresh

Source Database: appdb1  
 Owner: SYSMAN  
 Description: This is an example baseline.

▼ Scope Specification

Type: Table, Index, View, Trigger, Package, Package Body, Procedure, Function, Sequence, User, Role

Schemas to Include: HR, SCOTT  
 Grants for Users and Roles: Yes  
 Initialization Parameters: No

Versions

View View Changes Since Previous Version Generate DDL Delete

Select	Version	Objects	Creation Date	Job Status	DDL
<input type="radio"/>	1	106	Jan 7, 2009 1:27:24 AM (UTC-08:00)	Succeeded	Not Generated

Home | Targets | Deployments | Alerts | Compliance | Jobs | Reports | Setup | Preferences | Help | Logout

Copyright © 1996, 2009, Oracle and/or its affiliates. All rights reserved.  
 Oracle is a registered trademark of Oracle Corporation and/or its affiliates.  
 Other names may be trademarks of their respective owners.  
[About Oracle Enterprise Manager](#)

## Working With A Dictionary Baseline Version

Within a single dictionary baseline version, you can examine individual object attributes, generate DDL for individual objects, and generate DDL for all the objects in the baseline version. You cannot modify object definitions captured in baseline versions, since they are intended to represent the state of objects at a particular point in time.

- Viewing a baseline object displays the object's attributes graphically.
- Selecting a baseline object and specifying "Generate DDL" displays the DDL used to create the object. [Figure 7-3](#) shows the DDL generated for a baseline object.



**Figure 7–3 DDL Generated for a Selected Baseline Object**

ORACLE Enterprise Manager 10g  
Grid Control

Home Targets Deployments Alerts Compliance Jobs Reports

Hosts | Databases | Web Applications | Services | Systems | Groups | All Targets

Dictionary Baselines > Baseline: mv baseline > Baseline: mv baseline[1] >

Generate DDL: EMP OK

```
CREATE TABLE "SCOTT"."EMP"
(
  "EMPNO" NUMBER(4,0) NOT NULL ENABLE,
  "ENAME" VARCHAR2(10),
  "JOB" VARCHAR2(9),
  "MGR" NUMBER(4,0),
  "HIREDATE" DATE,
  "SAL" NUMBER(7,2),
  "COMM" NUMBER(7,2),
  "DEPTNO" NUMBER(2,0)
) PCTFREE 10 PCTUSED 40 INITRANS 1 MAXTRANS 255 NOCOMPRESS LOGGING
STORAGE(INITIAL 16384 NEXT 16384 MINEXTENTS 1 MAXEXTENTS 505
PCTINCREASE 50 FREELISTS 1 FREELIST GROUPS 1 BUFFER_POOL DEFAULT)
TABLESPACE "SYSTEM" ;
```

OK

Home | Targets | Deployments | Alerts | Compliance | Jobs | Reports | Setup | Preferences | Help | Logout

Copyright © 1996, 2009, Oracle and/or its affiliates. All rights reserved.  
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.  
Other names may be trademarks of their respective owners.  
[About Oracle Enterprise Manager](#)

- Selecting a baseline version and specifying "Generate DDL" generates the DDL for all objects in the baseline version. While an effort is made to create objects in the correct order (for example, creating tables before indexes), the resulting DDL cannot necessarily be executed on a database to create the objects in the baseline version. For example, if you capture all the schema objects contained in schema APPL1, then try to execute the baseline version DDL on a database that does not contain User APPL1, the generated DDL will fail to execute.

Baseline versions are also used with other Change Management Pack applications, such as Compare and Synchronize. You can compare a baseline version to a database (or to another baseline version). You can also use a baseline version as the source of object definitions in Synchronize, allowing you to re-create the definitions in another database.

## Working With Multiple Dictionary Baseline Versions

When a baseline contains more than one version, you can examine the differences between the versions.

- To see what has changed between a version and the version that precedes it, select the version and specify "View Changes Since Previous Version." The display shows which objects have changed since the previous version (Figure 7–4), which have been added or removed, and which are unchanged. Selecting an object that has changed displays the differences between the object in the two versions.

**Figure 7-4 Viewing Changes Since the Previous Version**

ORACLE Enterprise Manager 10g  
Grid Control

Home Targets Deployments Alerts Com

Hosts | Databases | Web Applications | Services | Systems | Groups | All Targets

Dictionary Baselines > Baseline: my baseline >

**Changes Since Previous Version: my baseline[2]**

Previous Version 1  
Source Database appdb1  
Creation Date Jan 8, 2009 12:33:21 AM (UTC-08:00)

**Objects Change Summary**

Unchanged	103	Changed	2	New	1	Removed	1
-----------	-----	---------	---	-----	---	---------	---

**Objects**

Object Type Schema Object Name

All Types   Go

Show Changed

Schema	Name	Type
SCOTT	DEPT	TABLE
SCOTT	EMP	TABLE

- To see how an individual object has changed over all the versions of the baseline, select the object and specify "View Version History." The display identifies the versions in which the object was initially captured, modified, or dropped as shown in Figure 7-5. From this display, you can compare the definitions of the object in any two baseline versions.

**Figure 7-5 Viewing Version History of a Selected Baseline Object**

ORACLE Enterprise Manager 10g  
Grid Control

Home Targets Deployments Alerts Compliance

Hosts | Databases | Web Applications | Services | Systems | Groups | All Targets

Dictionary Baselines > Baseline: my baseline > Baseline: my baseline[3] >

**Version History: EMP**

Object Type Table  
Schema SCOTT

Baseline my baseline  
Source appdb1

**Versions**

Capture Date	Baseline Version	Change Type
Jan 8, 2009 12:53:33 AM (UTC-08:00)	3	New Version
Jan 8, 2009 12:33:21 AM (UTC-08:00)	2	New Version
Jan 7, 2009 1:27:24 AM (UTC-08:00)	1	New Object

**Compare Object Versions**

Compare the definitions of this object that were first captured in:

Left: Baseline Version 1

Right: Baseline Version 3

## Dictionary Comparisons

A dictionary comparison identifies differences in database object definitions between a baseline or a database and a baseline or a database, or two schemas within a single database/baseline.

A comparison specification is defined by left and right sources, scope, and owner. The scope specification describes the names and types of database object definitions to be included in the comparison and the schemas that contain these object definitions.

Comparisons identify differences in any attribute value between objects of any type. Use comparisons to create multiple versions of a comparison. Each version has a unique version number and a comparison date. Use these versions to associate comparisons of database/schemas made over time.

### Examples:

Comparisons show differences between definitions in the original baseline for your application and those in your current database. After creating a new comparison version, it identifies the differences between the original definitions at the start of the development cycle, and those same definitions at the current time.

Use another comparison specification to compare definitions from your most recent baseline with those in your previous baseline. With each newly created version of this comparison using the comparison specification, that comparison version identifies the differences in the definitions since the previous baseline.

## Defining Dictionary Comparisons

A dictionary comparison definition consists of left and right sources for metadata definitions, the scope specification, an optional schema map, and comparison options. Once created, a dictionary comparison definition cannot be modified. This ensures that each version of the dictionary comparison reflects changes in the databases being compared and not in the comparison's definition.

### Dictionary Comparison Sources

Each comparison has a left source and a right source. The comparison matches object definitions from the left and right sources. A source can be a database or a baseline version.

- When the source is a database, object definitions are taken directly from the database at the time the comparison version is created.
- When the source is a baseline, object definitions are taken from the specified baseline version. Using a baseline version allows you to compare a database (or another baseline version) to the database as it existed at a previous point in time. For example, a baseline version might represent a stable point in the application development cycle, or a previous release of the application.

For baseline sources, there are various ways to specify the version to be used.

- If you want a specific baseline version to be used in all versions of the comparison, specify the baseline version number. This is appropriate for comparing a well-defined previous state of the database, such as a release, to its current state.
- You can also request that the latest or next-to-latest version be used in the comparison. If you specify "Latest," you can also request that the baseline version be captured before the comparison takes place. This option allows you to capture a baseline and compare it to the other source in a single operation. For example, every night, you can capture a baseline version of the current state of a

development database and compare it to the previous night's baseline, or to a fixed baseline representing a stable point in development.

### **Scope Specification**

The scope specification for a dictionary comparison identifies the objects to compare in the left and right sources. Creating a comparison scope specification is the same as creating a baseline scope specification, described in the "Dictionary Baselines" section. As with baselines, you can specify object types and schemas to compare, or individual objects to compare.

### **Schema Map**

Normally, schema objects in one source are compared to objects in the same schema in the other source. For example, table APPL1.T1 in the left source is compared to APPL1.T1 in the right source.

However, there may be cases where you want to compare objects in one schema to corresponding objects in a different schema. For example, assume that there are two schemas, DEV1 and DEV2, which contain the same set of objects. Different application developers work in DEV1 and DEV2. You can use the optional schema map feature to allow you to compare objects in DEV1 to objects with the same type and name in DEV2.

To add entries to the schema map, expand the "Mapped Objects" section of the comparison "Objects" page. You can create one or more pairs of mapped schemas. Each pair designates a left-side schema and a corresponding right-side schema.

When using a schema map, you can compare objects within a single database or baseline version. In the example above, DEV1 and DEV2 can be in the same database. You specify that database as both the left and right source, and supply the schema map to compare objects in DEV1 to those in DEV2.

### **Comparison Options**

You can select several options to determine how objects are compared. These options allow you to disregard differences that are not significant. The options include the following:

- "Ignore Tablespace" and "Ignore Physical Attributes" – These two options allow you to compare stored objects without regard to the tablespaces in which they are stored or the settings of their storage-related attributes. This is useful when you are comparing objects in databases having different size and storage configurations, and you are not interested in differences related to the storage of the objects.
- "Match Constraints By Definition" or "By Name" – If you are more interested in the definitions of table constraints – for example, the columns involved in primary or unique constraints – choose "Match Constraints By Definition." This causes constraints with the same definitions to match; their names appear as differences (unless you also choose "Ignore Name Differences"). If the names of constraints are meaningful, choose "Match Constraints By Name." With this choice, constraints with the same names will match and their definitions will appear as differences.
- "Partitioned Objects: Ignore High Values" – Tables that are otherwise the same might have different partition high values in different environments. Choose this option to ignore differences in high values.
- "Logical SQL Compare" – Choose this option to ignore meaningless formatting differences in source objects such packages, package bodies, procedures and functions.

- "Compare Statistics" – Choose this option to compare optimizer statistics for tables and indexes.
- "Ignore Table Column Position" – Choose this option if tables that differ only in column position should be considered equal.

### **Creating Comparison Versions**

When you have finished defining the comparison, you specify when to create the first comparison version. You can create the first version immediately, or at a later time. You can also schedule new comparison versions at regular intervals.

In addition to scheduling comparison versions, you can create a new comparison version at any time by selecting the comparison and specifying "Repeat Now."

Comparisons processed after the initial version generally complete substantially faster than the initial comparison. Only those objects that have changed on the left or right side are compared in the new version. This also means that the storage required for additional comparison versions is only slightly larger than the storage used by the initial version, assuming a small percentage of objects has changed.

## **Working with Dictionary Comparison Versions**

A dictionary comparison version records the results of comparing the left and right sources in accordance with the scope specification. Objects in a comparison version have one of four states:

- Left Only – The object is present only in the left source.
- Right Only – The object is present only in the right source.
- Identical – The object is present in both left and right sources, and is the same.
- Not Identical – The object is present in both left and right sources, and is different.

**Figure 7-6 Comparison Version Page Showing a List Of Not Identical Objects**

The screenshot shows the Oracle Enterprise Manager 10g interface for a dictionary comparison. The page title is "Comparison: my comparison[1]". It displays the following information:

- Left Source:** appdb1, Left Type: Database, Owner: SYSMAN
- Right Source:** database\_2, Right Type: Database, Creation Date: Jan 8, 2009 1:23:55 AM (UTC-08:00)
- Initialization Parameters:** Identical: 325, Not Identical: 16, Left Only: 0, Right Only: 0
- Objects:** A search filter is set to "Not Identical - Present in both sources but attributes are not identical". A table lists 10 objects that are not identical.

Select	Schema	Name	Result	Type	Most Recent Comment	Details
<input type="checkbox"/>	OE	CUSTOMERS		TABLE		
<input type="checkbox"/>	HR	DEPARTMENTS		TABLE		
<input type="checkbox"/>	SCOTT	DEPT		TABLE		
<input type="checkbox"/>	SCOTT	EMP		TABLE	Here I can enter a comment such as why this particular table is different.	
<input type="checkbox"/>	HR	EMPLOYEES		TABLE		
<input type="checkbox"/>	HR	JOB_HISTORY		TABLE		
<input type="checkbox"/>	OE	ORDERS		TABLE		
<input type="checkbox"/>		SYSTEM		TABLESPACE		
<input type="checkbox"/>		TEMP		TABLESPACE		

The page lists all versions of a comparison shows the number of objects in each state within each version. On the Comparison version page shown in Figure 7-6, you can see the objects in each state individually. Objects that are "Not Identical" can be selected to view the differences, and to generate DDL for the left and right definitions as shown in Figure 7-7.

**Figure 7–7 Viewing Differences in the Attributes of a Selected Not Identical Object**

ORACLE Enterprise Manager 10g  
Grid Control

Home Targets Deployments Alerts Compliance Jobs Reports

Hosts | Databases | Web Applications | Services | Systems | Groups | All Targets

Dictionary Comparisons > Comparison: my comparison > Comparison: my comparison[1] >

View Differences: EMP Return

Left Source **appdb1** Right Source **database\_2**  
Object Type **Table** Schema **SCOTT**  
Object Name **EMP**

**Differences**

[Expand All](#) | [Collapse All](#)

Attributes	appdb1	database_2
Table		
Columns		
EMAIL_ADDRESS		Not Present
Name	EMAIL_ADDRESS	
Data Type	VARCHAR2	
Column Position	9	
Length	100	
Constraints		
Unique Keys		
UK_EMAIL_ADDRESS		Not Present

TIP To view the DDL for the left or right object, click on 'Left DDL' or 'Right DDL' below.

[View Comments](#)

[Left DDL](#) [Right DDL](#)

You can take two further actions to record information about objects in a comparison version:

- You can add a comment to an object. For example, the comment might explain why two objects are different.
- You can ignore the object. Ignoring the object removes it from lists of comparison version objects. You might ignore an object that is different if you decide that the difference is not important.

## Dictionary Synchronizations

A dictionary synchronization synchronizes differences in database object definitions between two databases or a baseline and a database. The basic action of a database synchronization is to create or modify selected object definitions in a database to match those in another database or in a baseline.

Synchronizations are generated using synchronization specifications. For synchronizations, the scope specification does not include the names of individual objects. It can only specify the types, and the schemas to be included or excluded. You can additionally supply a prefix to limit the objects selected to those whose names start with that prefix.

Dictionary synchronizations synchronize differences in any attribute value between objects of any supported type. Use synchronization specifications to create multiple versions of a synchronization. Each version has a unique version number and a synchronization date. Use these versions to associate synchronizations of database/schemas made over time.

## Defining Dictionary Synchronizations

### Source and Destination

The synchronization source provides the object definitions (and optionally, the data) from which the destination database is synchronized. A synchronization source may be a database, or a baseline version. If the source is a baseline version, it is not possible to propagate data to the destination, since a baseline does not capture data.

The synchronization destination must always be a database. The purpose of synchronization is to create or modify object definitions in the destination to match those in the source.

The options for specifying which version of a source baseline to use are similar to those used with dictionary comparisons. You can specify a fixed baseline version, or "Latest" or "Latest-1". If you specify "Latest," you can also request that the baseline version be captured first before synchronizing the destination from it.

When defining a baseline to be used as the source for synchronization, it is important that the baseline contain all the objects to be synchronized. For this reason, the baseline's scope specification should include at least all the schemas and object types that will be synchronized. The baseline should also include User and Role objects, along with privilege grants. A baseline to database synchronization is recommended in environments where changes are expected to be applied to the source database frequently thus necessitating the need for a point-in-time snapshot of the database, i.e. a baseline as the source of the synchronization.

### Scope Specification

Defining a scope specification for a dictionary synchronization is similar to defining a scope specification for a dictionary baseline or comparison. However, there are restrictions on what you can include in the scope specification.

- You cannot specify individual objects for synchronization. You must specify object types and either schemas to include or schemas to exclude.
- Certain schemas, such as SYS and SYSTEM, cannot be synchronized.
- You cannot directly include User and Role objects for synchronization. (However, Users and Roles are automatically included as needed during the synchronization process.)
- The following object types must be selected as a group: Table, Index, Cluster, Materialized View, and Materialized View Log.

The scope specification for a synchronization should be carefully tailored to the application. Do not include more schemas than are needed. For example, if you are synchronizing a module of a large application, include only those schemas that make up the module. Do not attempt to synchronize a large application (or the entire database) at one time.

### Schema Map

The definition and use of the schema map is the same in dictionary synchronizations as in dictionary comparisons. When you use a schema map, object definitions in each schema map are synchronized to the mapped schema in the destination, rather than to the schema with the same name. In addition, schema-qualified references (other than those contained in PL/SQL blocks and view queries) are changed according to the schema map.

For example, assume the schema map has two entries, as follows:



- DEV1A -> DEV2A
- DEV1B -> DEV2B

Table DEV1A.T1 has an index, DEV1A.T1\_IDX. It also has a foreign key constraint that refers to DEV1B.T2. Synchronize will create objects as follows:

- Table DEV2B.T2
- Table DEV2A.T1, with a foreign key reference to table DEV2B.T2
- Index DEV2A.T1\_IDX, on table DEV2A.T1

### Synchronization Options

Dictionary synchronization options are similar to the options you can specify with dictionary comparisons. In synchronization, the options perform two functions:

- During initial comparison of source and destination objects, the options determine whether differences are considered meaningful. For example, if the "Ignore Tablespace" option is selected, tablespace differences are ignored. If two tables are identical except for their tablespaces, no modification to the destination table will occur.
- When generating the script that creates objects at the destination, some options control the content of the script. For example, if "Ignore Tablespace" is selected, no TABLESPACE clauses are generated.

In addition to the options provided with dictionary comparison, the following options are specific to Synchronize:

- "Preserve Data In Destination" and "Copy Data From Source" – These two options control whether table data is copied from the source database to the destination database. (The option is not available if the source is a baseline.) By default, Synchronize preserves data in destination tables. Choosing "Copy Data From Source" causes Synchronize to replace the destination data with data from the source table.
- "Drop Destination-Only Schema Objects" – Choosing this option directs Synchronize to drop schema objects that are present in the destination but not the source. For example, if the source contains table DEV1.T1 and the destination contains DEV1.T1 and DEV1.T2, Synchronize will drop DEV1.T2 if this option is chosen. This action applies only to schema objects that are within the scope specification. By default, Synchronize does not drop destination-only objects. Synchronize never drops non-schema objects.

### Synchronization Mode

The next step in defining a synchronization is to choose the synchronization mode. There are two options:

- Unattended synchronization mode carries out the entire synchronization process in one step, ending with execution of the synchronization script. However, if an error is detected that makes it impossible to generate a correct script, the process will terminate without attempting to execute the script.
- Interactive synchronization mode pauses after initial comparison of the source and destination objects, and again after generation of the synchronization script. Interactive mode gives you a chance to examine the results of comparison and script generation, and to take appropriate action before proceeding with the next step.

### Creating Synchronization Versions

When you have finished defining the synchronization, you specify when to create the first synchronization version. You can create the first version immediately, or at a later time. You can also schedule new synchronization versions at regular intervals.

Depending on the synchronization mode you select, the synchronization may run to completion (unattended mode), or pause after initial object comparison (interactive mode). In the latter case, you run each subsequent phase of the synchronization in a new job.

When using interactive mode, the destination database should not be modified from the time the objects are initially compared until the synchronization script has executed. Otherwise, the script may encounter problems. For example, assume the source has a table that the destination does not have. Object comparison notes the source-only table, and the generated script includes statements to create the table. However, after object comparison but before script execution, you manually create the table at the destination database. The script will fail when it attempts to create the table because the table already exists.

Figure 7–8 shows a list of synchronizations.

**Figure 7–8 List of Synchronizations**

**Dictionary Synchronizations**

A Dictionary Synchronization operation makes the definitions of set of objects in the destination database the same as those in the source. The source can be either a database or a baseline version. You can define what you want to synchronize once using the Objects step, and synchronize again at later times. You can copy data from the source database to destination. Synchronization script will be generated and can be executed immediately or at a later time.

Search:  Name

Select	Name	Source	Destination	Versions	Owner	Most Recent Version	Most Recent Status	Pending Action
<input checked="" type="radio"/>	my synchronization	appdb1	database_2	1	SYSMAN	Jan 9, 2009 1:34:04 AM (UTC-08:00)	✓ Comparison Succeeded	<a href="#">Generate Script</a>
<input type="radio"/>	reset dict changes	my baseline[1]	database_2	2	SYSMAN	Jan 9, 2009 1:32:32 AM (UTC-08:00)	✓ Execution Succeeded	
<input type="radio"/>	apply dict changes	my baseline [Latest]	appdb1	1	SYSMAN	Jan 9, 2009 1:03:57 AM (UTC-08:00)	✓ Execution Succeeded	
<input type="radio"/>	my incorrect sync	appdb1	database_2	1	SYSMAN	Jan 8, 2009 9:59:18 AM (UTC-08:00)	✗ Generated With Errors	

**Related Links**  
[Dictionary Baselines](#)      [Dictionary Comparisons](#)

For more information about the process of synchronization, see [Working with Dictionary Synchronization Versions](#).

### Creating a Synchronization Definition from a Comparison

You can use a dictionary comparison as the starting point for synchronization. Select the comparison, then choose "Synchronize." This creates a new synchronization with the following initial information from the comparison:

- Source and destination, from the comparison’s left and right sources, respectively. This means that you cannot create a synchronization from a comparison whose right source is a baseline.

- Scope specification. Note that some comparison scope specification options are not available in a synchronization. For example, you cannot synchronize individual schema objects, User objects, or Role objects.
- Comparison options

## Working with Dictionary Synchronization Versions

Each synchronization version represents an attempt to modify the destination objects selected by the scope specification to match the corresponding source objects. (It is "an attempt" because the process may not complete, for various reasons.) This section describes how a dictionary synchronization version is processed, and how you can monitor and control the process.

### The Dictionary Synchronization Cycle

There are three steps involved in processing a synchronization version. As noted previously, you can combine these steps into one (by choosing "Unattended Mode") or run each step separately (by choosing "Interactive Mode"). In either case, all three steps must be carried out when processing a successful synchronization version.

The following sections detail each of the three steps and describe what you can do following each step, when operating in interactive mode.

#### Object Comparison Step

The first step is to compare objects in the source to corresponding objects in the destination. Only those objects selected by the scope specification are compared. At the end of this step, the synchronization version has recorded all the objects. Each object is in one of the following states:

- Source Only
- Destination Only
- Identical
- Not Identical

In interactive mode, you can view the objects that are in each state, or all objects at once. For objects that are not identical, you can view the differences between the objects. At this stage, you can anticipate what will happen to each destination object:

- Source-only objects will be created in the destination.
- Destination-only objects will be unaffected, unless you chose the "Drop Destination-Only Schema Objects" option, in which case they will be dropped from the destination.
- Identical objects will be unaffected. However, if you chose the "Copy Data From Source" option, the data in tables that are identical will be replaced with data from the source.
- Non-identical objects will be modified to match the source object. Depending on the differences, the modification may be done with ALTER statements or by dropping and re-creating the object.

Before proceeding to script generation ([Figure 7-9](#)), you can exclude objects from the synchronization. For example, if you notice a source-only view that you do not want to create at the destination, you can exclude it now.

**Figure 7–9 Continue to Script Generation Step in the Interactive Mode**

ORACLE Enterprise Manager 10g  
Grid Control

Home Targets Deployments Alerts Compliance Jobs Reports

Hosts | Databases | Web Applications | Services | Systems | Groups | All Targets

**Synchronize: Continue To Script Generation** Cancel Continue

Synchronization Name: my synchronization      Version: 1  
 Source Database: appdb1      Creation Date: Jan 9, 2009 1:34:04 AM (UTC-08:00)  
 Destination Database: database\_2      Job Status: Comparison Succeeded  
 Owner: SYSMAN  
 Description: Synchronize database\_2 with changes from appdb1.

Use this page to view comparison results and selectively exclude objects from the synchronization operation. The excluded objects will be ignored during script generation.

Comparison Summary    Source Only > 2    Destination Only < 1    Identical = 38    Not Identical 2

Object Type: All Types    Schema:    Object Name:    Go

Show: All Differences - Source Only, Destination Only and Non-Identical objects    Results

Exclude From Synchronization    Undo Exclude

Select All | Select None

Select	Schema	Name	Result	Type
<input type="checkbox"/>	SCOTT	BONUS	<	TABLE
<input type="checkbox"/>	SCOTT	DEPT	2	TABLE
<input type="checkbox"/>	SCOTT	EMP	2	TABLE
<input type="checkbox"/>	SCOTT	LOCATION	>	TABLE
<input type="checkbox"/>	SCOTT	UK_EMAIL_ADDRESS	>	INDEX

### Script Generation Step

During script generation, Synchronize uses the results of the comparison step to create the PL/SQL script that will create and modify objects at the destination so that it matches the source. As part of script generation, several activities take place:

- Dependency analysis assures that the destination database provides a suitable environment for each object and that objects are created and modified in the correct order.
- Change analysis determines whether an object can be modified with ALTER statements or if it must be dropped and re-created.
- Messages are placed in the impact report for the synchronization version. The messages provide information about the synchronization process, and may report one or more error conditions that make it impossible to generate a usable script.
- The DDL statements needed to carry out the synchronization are generated and recorded in the synchronization version.

Dependency analysis examines each object to determine its relationships with other objects. It makes sure these other objects exist (or will be created) in the destination. Some examples of these relationships are:

- A schema object depends on the User object that owns it. For example, table DEV1.T1 depends on user DEV1.
- An index depends on the table that it is on. For example, index DEV1.T1\_IDX depends on table DEV1.T1.
- A table that has a foreign key constraint depends on the table to which the constraint refers.
- A source object such as a package body depends on other source objects, tables, views, and so on.

- A stored object depends on the tablespace in which it is stored, unless you choose "Ignore Tablespace."

The relationships established during dependency analysis are used later in the script generation process to make sure script statements are in the correct order.

Dependency analysis may determine that a required object does not exist in the destination database. For example, a schema object's owner may not exist in the destination database. Or, a table may have a foreign key constraint on another table that is in a different schema. There are several possible outcomes.

- If the required object is in the source and is selected by the scope specification, Synchronize creates the object. For example, if DEV1.T1 has a foreign key constraint that refers to DEV2.T2, and both DEV1 and DEV2 are in the scope specification, Synchronize creates DEV2.T2.
- If the required object is a user or role, and the object is available in the source, Synchronize automatically includes the user or role and creates it at the destination. This occurs even though User and Role objects are not part of the Synchronize scope specification.
- If a required schema object is in the source but is not selected by the scope specification, Synchronize does not automatically include the object; instead, it places an Error-level message in the impact report. This restriction prevents uncontrolled synchronization of objects outside the scope specification. It is for this reason that scope specifications should include all the schemas that make up the application or module.
- If the source is a baseline version, it may not include the required object. For example, a baseline might not capture Users and Roles. Synchronize cannot look for objects outside the baseline version, so it places an Error-level message in the impact report as shown in [Figure 7-10](#). This is why it is important to include Users, Roles, and privilege grants in any baseline that will be used for synchronization.

**Figure 7–10 Error-level Impact Report Messages Resulting from a Required Schema Object Not Selected by the Scope Specification**

**View Synchronization: my incorrect sync[1]**

Source Database	appdb1	Creation Date	Jan 8, 2009 9:59:18 AM (UTC-08:00)
Destination Database	database_2	Job Status	Generated With Errors
Owner	SYSMAN	Most Severe Message	✘ Error
Description	EXAMPLE of an incorrect sync attempt where an in-scope view DEV2.V2 refers to an an out-of-scope table DEV3.T3. To correct this error, a new sync needs to be created that includes the DEV3 schema as well.		

Message Summary: Error ✘ 2, Warning 0, Information i 3

Select	Severity	Object Type	Schema	Object Name	Message	Recommended Action
<input checked="" type="radio"/>	✘	TABLE	DEV3	T3	An object that is being synchronized depends on this object, but this object cannot be included because it is out of scope.	Create a new synchronization that includes schema DEV3 and type TABLE so that this table is also synchronized.
<input type="radio"/>	✘	VIEW	DEV2	V2	This view refers to TABLE DEV3.T3, which was not found in the destination or source.	Create a new synchronization that includes table DEV3.T3.
<input type="radio"/>	i	USER		DEV1	This object was included from the source. TABLE DEV1.T1, and possibly other objects, depend on this object.	Make sure that it is appropriate to include this object.
<input type="radio"/>	i	USER		DEV2	This object was included from the source. USER DEV1, and possibly other objects, depend on this object.	Make sure that it is appropriate to include this object.
<input type="radio"/>	i	USER		DEV3	This object was included from the source. USER DEV2, and possibly other objects, depend on this object.	Make sure that it is appropriate to include this object.

At the end of the script generation step, Synchronize has added the impact report and the script to the synchronization version. In interactive mode, you can examine the script and impact report before proceeding to script execution.

The impact report contains messages about situations that were encountered during script generation. There are three types of messages:

- Informational messages require no action on your part. They simply report something of interest. For example, when script generation automatically includes a User object, it adds an informational message to the impact report.
- Warning messages report a situation that requires your attention, but that may not require action. For example, if Synchronize is unable to determine if a reference is a source object can be resolved, it adds a warning message to the impact report. You need to verify that situations reported in warning messages will not prevent script execution.
- Error messages indicate a situation that will prevent script execution if not corrected. For example, if Synchronize is unable to locate a required dependency object, it adds an error message to the impact report. Depending on the message, you may be required to create a new synchronization. For example, if the dependency object is not in the synchronization scope, or if the source is a baseline that does not contain the dependency object, you will need to create a new synchronization with an expanded scope or a different source baseline. In other cases, you can resolve the situation by excluding one or more objects from the synchronization and regenerating the script.

The script display contains the statements of the generated script in the order they will be executed. You can examine the script if you have any concerns about its correctness.

The display allows you to locate statements that are associated with a particular object or object type.

**Figure 7–11 Continuing To Script Execution Step in the Interactive Mode (Impact Report Tab)**

ORACLE Enterprise Manager 10g

Grid Control

Home Targets Deployments Alerts Compliance Jobs Reports

Hosts | Databases | Web Applications | Services | Systems | Groups | All Targets

**Synchronize: Continue To Script Execution** Cancel Continue

Synchronization Name **my synchronization** Version **1**  
 Source Database **appdb1** Creation Date **Jan 9, 2009 1:52:59 AM (UTC-08:00)**  
 Destination Database **database\_2** Job Status **Generated With Warnings**  
 Owner **SYSMAN** Most Severe Message **Warning**  
 Description **synchronize database\_2 with changes from appdb1.**

**Impact Report** [Script](#)

Review Impact Report messages for the synchronization operation. Take suggested user actions to correct any errors.

Message Summary Error 0 Warning **2** Information 0

Object Type Schema Object Name  
 All Types   Go

Severity	Object Type	Schema	Object Name	Message	Recommended Action
Warning	TABLE	SCOTT	DEPT	Data cannot be preserved appropriately for this table since constraint DEPT_LOC_FK added	Ensure the data values do not violate the constraints
Warning	TABLE	SCOTT	EMP	Data cannot be preserved appropriately for this table since constraint UK_EMAIL_ADDRESS added	Ensure the data values do not violate the constraints

[About Oracle Enterprise Manager](#)

Following script generation, you can continue to script execution unless an error was encountered during the script generation step (Figure 7–11 and Figure 7–12). In this case the impact report will contain one or more Error-level messages detailing the problem and solution. In some cases, you may be able to solve the problem by selecting "Regenerate Script," excluding an object from the synchronization, and regenerating the script.

**Figure 7–12 Continuing To Script Execution Step in the Interactive Mode (Script Tab)**

**Synchronize: Continue To Script Execution** Cancel Continue

Synchronization Name: my synchronization      Version: 1  
 Source Database: appdb1      Creation Date: Jan 9, 2009 1:52:59 AM (UTC-08:00)  
 Destination Database: database\_2      Job Status: Generated With Warnings  
 Owner: SYSMAN      Most Severe Message: Warning  
 Description: synchronize database\_2 with changes from appdb1.

[Impact Report](#)   [Script](#)   [Save Full Script](#)

**Search**

Script Line Type: All Types      Object Type: All Types  
 Schema:       Object Name:

Script Line Type	Object Type	Schema	Object Name	Script Line	Details
Comment				-- Script Execution Controller	<a href="#">Details</a>
SQL	TABLE	SCOTT	BONUS	DROP TABLE "SCOTT"."BONUS" PURGE	<a href="#">Details</a>
SQL	TABLE	SCOTT	LOCATION	CREATE TABLE "SCOTT"."LOCATION" ("LOCATION_ID" NUMBER(4,0), "STREET_ADDRESS" VARCHAR...	<a href="#">Details</a>
SQL	TABLE	SCOTT	LOCATION	ALTER TABLE "SCOTT"."LOCATION" ADD PRIMARY KEY ("LOCATION_ID") ENABLE	<a href="#">Details</a>
SQL	TABLE	SCOTT	DEPT	ALTER TABLE "SCOTT"."DEPT" DROP ("LOC")	<a href="#">Details</a>
SQL	TABLE	SCOTT	DEPT	ALTER TABLE "SCOTT"."DEPT" ADD ("LOCATION_ID" NUMBER(4,0))	<a href="#">Details</a>
SQL	TABLE	SCOTT	DEPT	ALTER TABLE "SCOTT"."DEPT" ADD CONSTRAINT "DEPT_LOC_FK" FOREIGN KEY ("LOCATION_ID") REFERE...	<a href="#">Details</a>
SQL	TABLE	SCOTT	EMP	ALTER TABLE "SCOTT"."EMP" ADD ("EMAIL_ADDRESS" VARCHAR2(100))	<a href="#">Details</a>
SQL	INDEX	SCOTT	UK_EMAIL_ADDRESS	CREATE UNIQUE INDEX "SCOTT"."UK_EMAIL_ADDRESS" ON "SCOTT"."EMP" ("EMAIL_ADDRESS")	<a href="#">Details</a>
SQL	TABLE	SCOTT	EMP	ALTER TABLE "SCOTT"."EMP" ADD CONSTRAINT "UK_EMAIL_ADDRESS" UNIQUE ("EMAIL_ADDRESS") ENABL...	<a href="#">Details</a>

There may be cases where you need to create a new version of the synchronization in order to correct the problem. For example, if you need to modify the definition of an object in the source or destination or add an object in the destination, you will need to create a new version. This allows the new or modified object to be detected during the comparison step. In this case, the old version becomes "abandoned" since you cannot continue to script generation.

**Script Execution Step**

Following successful script generation, the script is ready to execute. In unattended mode, the script executes as soon as script generation completes. In interactive mode, you proceed to script execution as soon as you have reviewed the impact report and the script.

The script executes in the Enterprise Manager job system. [Figure 7–13](#) shows the output of the Synchronization job. Once script execution is complete, you can view the execution log. If the script fails to execute successfully, you may be able to correct the problem and re-start the script execution from the point of failure. For example, if the script fails due to lack of space in a tablespace, you can increase the size of the tablespace, then re-start the script.



Figure 7–13 Synchronization Script Execution Step Output of Enterprise Manager Job

The screenshot displays the Oracle Enterprise Manager 10g interface. At the top, there are navigation tabs for 'Home', 'Targets', 'Deployments', and 'Alerts'. Below these, the 'Job Activity | Job Library' section is active, showing 'Job Activity > Execution: 2 >'. The main heading is 'Step: Synchronization Script Execution'. A 'Page Refre' link is visible on the right.

The execution details are as follows:

Status	Succeeded	Started	Jan 9, 2009 2:09:37 AM (UTC-08:00)
Exit Code	0	Ended	Jan 9, 2009 2:09:40 AM (UTC-08:00)
Step ID	906	Step	
Targets	database_2	Elapsed Time	3 seconds
		Management Service	stapn32.us.oracle.com:7664_Management_Service

A note indicates: [TIP Management Service from which the job step was dispatched.](#)

The 'Output Log' section contains the following text:

```
SQL*Plus: Release 11.1.0.7.0 - Production on Fri Jan 9 02:09:38 2009
Copyright (c) 1982, 2008, Oracle. All rights reserved.

SQL> Connected.
SQL> -- Script Execution Controller
Starting Synchronization
USER is "SYSTEM"
DROP TABLE "SCOTT"."BONUS" PURGE
CREATE TABLE "SCOTT"."LOCATION"
(
  "LOCATION_ID"
    NUMBER(4,0),
  "STREET_ADDRESS" VARCHAR2(40),
  "CITY" VARCHAR2(30) CONSTRAINT
"LOC_CITY_NN" NOT NULL ENABLE,
  "STATE_PROVINCE" VARCHAR2(25),
  "POSTAL_CODE"
    VARCHAR2(12)
)
ALTER TABLE "SCOTT"."LOCATION" ADD PRIMARY KEY ("LOCATION_ID") ENABLE
ALTER TABLE "SCOTT"."DEPT" DROP ("LOC")
ALTER TABLE "SCOTT"."DEPT" ADD ("LOCATION_ID" NUMBER(4,0))
ALTER TABLE "SCOTT"."DEPT" ADD CONSTRAINT "DEPT_LOC_FK" FOREIGN KEY
("LOCATION_ID") REFERENCES "SCOTT"."LOCATION"("LOCATION_ID") ENABLE
ALTER TABLE "SCOTT"."EMP" ADD ("EMAIL_ADDRESS" VARCHAR2(100))
CREATE UNIQUE INDEX "SCOTT"."UK_EMAIL_ADDRESS" ON "SCOTT"."EMP"
("EMAIL_ADDRESS")
ALTER TABLE "SCOTT"."EMP" ADD CONSTRAINT "UK_EMAIL_ADDRESS" UNIQUE
("EMAIL_ADDRESS") ENABLE
Completed Synchronization. Starting cleanup phase.
Starting cleanup of recovery tables
Completed cleanup of recovery tables
Starting cleanup of generated procedures
Completed cleanup of generated procedures
Script execution complete
SQL> set ver on
SQL> Disconnected from Oracle Database 11g Enterprise Edition Release 11.1.0.7.0 - Production
With the Partitioning, Data Mining and Real Application Testing options
```

## Creating Additional Synchronization Versions

Following processing of the initial synchronization version, you can create additional versions of the synchronization. Select the synchronization, then choose "Synchronize Again." Note that you cannot choose a different source or destination, or modify the scope specification or synchronization options, when creating a new synchronization version. However, you can choose a different mode (Unattended or Interactive) when starting a new synchronization version.

Creating additional synchronization versions allows you to propagate incremental changes in the source to the destination. For example, you can refresh a test database from the latest changes made to a development system.

Synchronizations processed after the initial version generally complete substantially faster than the initial synchronization.

---

---

# Middleware Management

This chapter describes how you can use Enterprise Manager to manage middleware targets such as Oracle WebLogic Server, Oracle Application Server, Oracle BPEL Process Manager, and so on.

This chapter covers the following:

- [Introduction to Middleware Management](#)
- [Monitoring Middleware Targets](#)
- [Administering Middleware Targets](#)
- [Maintaining Middleware Targets](#)
- [Managing Identity Management Targets](#)

## Introduction to Middleware Management

Middleware essentially refers to the middle tier of your architecture that provides the runtime infrastructure for Web applications. Middleware consists of targets such as Oracle WebLogic Server, Oracle Application Server, Oracle BPEL Process Manager, Oracle BI Suite Enterprise Edition, and so on. While each middleware target has its own console to help you monitor that particular target, Grid Control offers a single-window solution wherein you can discover and centrally monitor many instances of a middleware target and also many middleware targets in one console.

For example, an Oracle Application Server can be managed using a standalone Application Server Control that comes with every Oracle Application Server. Alternatively, it can be managed using Grid Control along with other Oracle Application Servers and a host of other middleware targets.

This section introduces you to Application Server Control and Grid Control, and explains under what circumstances they come into use.

## Out-of-Box Management Using Oracle Application Server Control

Every Oracle Application Server instance is installed with an Application Server Control to help you manage that particular application server instance. Application Server Control provides Web-based management features designed to monitor and administer application server instances, farms, and clusters. You can also deploy applications, monitor real-time performance, manage security, and configure the application server components.

Application Server Control relies heavily on various underlying technologies to discover, monitor, and administer the application servers in an environment. It consists of the Application Server Control console and its underlying technologies:

- Oracle Dynamic Monitoring Service (DMS)
- Oracle Process Management Notification (OPMN)
- Distributed Configuration Management (DCM)
- A local version of the Oracle Management Agent specifically designed to gather monitoring data

However, you can manage only one application server at a time using Application Server Control. Typically, you have multiple application servers in your enterprise configuration, and managing all these instances using individual Application Server Controls becomes very difficult.

For centralized management and additional management functionality (for example, application service level management, deployments, historical data collections for performance trending alerts, and so on), you can use Grid Control.

## Centralized Management Using Enterprise Manager

While Application Server Control provides standalone management for a single application server instance and its components, Grid Control provides centralized management of multiple application servers in your environment.

For example, if you have 10 application servers installed on ten different hosts, then you can manage all these ten application servers and hosts through a single window using Grid Control. With the help of Management Agents deployed on each host, Grid Control automatically discovers the application servers on these hosts and begins monitoring them using default monitoring levels, notification rules, and other default settings. Similarly, if you have multiple logical entities such as Oracle WebLogic Server Domains across your network, you can discover all of them in Grid Control and monitor them using a single console.

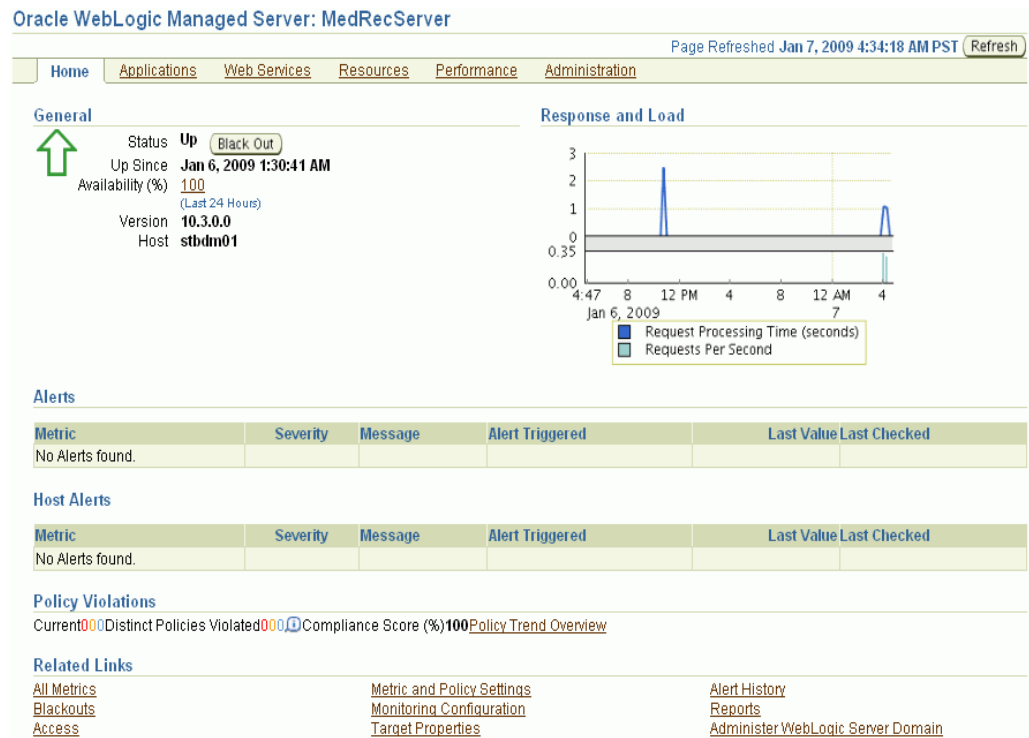
Grid Control provides its own application server home page that provides easy access to key information required by the administrators. Grid Control also provides a separate home page for logical entities such as Oracle WebLogic Server Domains and Oracle WebLogic Server Cluster.

At a high level, the home page in Grid Control ([Figure 8-1](#)) provides:

- Application server status, responsiveness, and performance data
- Resource usage for the application server and its components
- Functionality to start, stop, and restart any of those core components
- Alerts and diagnostic drill-downs so you can identify and resolve problems quickly
- Metrics for more detailed information about the health of the application server
- Links to other pages in Grid Control that might be helpful in accomplishing your given task

The following is the Oracle WebLogic Managed Server Home page:

Figure 8-1 Oracle WebLogic Managed Server Home Page



## Monitoring Middleware Targets

Grid Control provides a comprehensive set of features for monitoring middleware targets in your environment. You can view a summary of the most critical information pertaining to your middleware targets. You can also monitor their performance, view the alerts and policy violations that were generated, track the configuration changes that were made over a period of time, and perform other administrative tasks on them.

## What Gets Monitored?

Grid Control helps you monitor the following middleware targets:

**Table 8–1** *Middleware Targets Monitored*

Monitored Target	Services Offered by Grid Control	Related Links
Oracle WebLogic Server Domain, Oracle WebLogic Server Cluster, and Oracle WebLogic Managed Server	<ul style="list-style-type: none"> <li>■ Helps you to manually discover Oracle WebLogic Server Domains and monitor their overall health. As part of this discovery, the associated Oracle WebLogic Server Clusters and their Oracle WebLogic Managed Servers are also discovered.</li> <li>■ Allows you to diagnose, notify, and correct performance and availability problems in these servers.</li> <li>■ Helps you to monitor the web services and resources used by these servers.</li> <li>■ Helps you to monitor the applications running on these Oracle WebLogic Managed Servers.</li> <li>■ Helps you to monitor the web modules and EJB modules of these applications running on the servers.</li> <li>■ Allows you to create jobs to automatically start, stop, or restart the servers.</li> <li>■ Allows you to automatically create a service dependency for the servers.</li> </ul>	
Oracle Application Server	<ul style="list-style-type: none"> <li>■ Automatically discovers Oracle Application Servers in your environment and monitors their status, responsiveness, and performance.</li> <li>■ Helps you to view the J2EE and Web applications deployed on the application servers.</li> <li>■ Helps you to perform administrative tasks like patching or cloning Oracle homes, backing up or recovering application servers, and viewing, comparing, or searching configuration information</li> <li>■ Also helps you to view the topology.</li> </ul>	<a href="#">Monitoring Performance</a> <a href="#">Viewing Topology</a> <a href="#">Creating Services from J2EE or Web Applications</a> <a href="#">Diagnosing Performance Issues with Top Reports</a> <a href="#">Managing Backup and Recovery of the Application Server Environment</a> <a href="#">Cloning and Patching the Middleware Environment</a>
Oracle HTTP Server	<ul style="list-style-type: none"> <li>■ Automatically discovers and monitors Oracle HTTP Servers running within the application servers.</li> <li>■ Provides a host of metrics to gauge the server performance and virtual host performance.</li> <li>■ Helps you to view the top URLs being accessed.</li> <li>■ Also helps you to perform enterprise configuration management tasks like viewing, comparing, and searching configuration information.</li> </ul>	<a href="#">Diagnosing Performance Issues with Top Reports</a> <a href="#">Monitoring Performance</a> <a href="#">Managing Configurations</a>

**Table 8–1 (Cont.) Middleware Targets Monitored**

Monitored Target	Services Offered by Grid Control	Related Links
OC4J	<ul style="list-style-type: none"> <li>▪ Automatically discovers and monitors OC4J instances running within the application servers.</li> <li>▪ Helps you to view the active applications running on them.</li> <li>▪ Also helps you to perform enterprise configuration management tasks like viewing, comparing, and searching configuration information.</li> </ul>	<a href="#">Diagnosing Performance Issues with Top Reports</a> <a href="#">Managing Configurations</a>
OracleAS Web Cache	<ul style="list-style-type: none"> <li>▪ Automatically discovers and monitors OracleAS Web Cache instances running within the application servers.</li> <li>▪ Helps you to view the metrics associated with this target to analyze their performance.</li> <li>▪ Also help you to perform enterprise configuration management tasks like viewing, comparing, and searching configuration information.</li> </ul>	<a href="#">Monitoring Performance</a> <a href="#">Managing Configurations</a>
Oracle BPEL Process Managers	<ul style="list-style-type: none"> <li>▪ Automatically discovers and monitors Oracle BPEL Process Managers deployed on the application servers.</li> <li>▪ Helps you to view a list of domains and processes for a BPEL target, and the partner links associated with a particular BPEL process.</li> <li>▪ Helps you to create infrastructure services for BPEL targets, aggregate services for their processes, and SOAP Tests for the partner links.</li> <li>▪ Helps you to run deployment procedures for provisioning BPEL processes on BPEL Process Managers.</li> <li>▪ Helps you to perform enterprise configuration management tasks like viewing, comparing, and searching configuration information.</li> <li>▪ As part of the configuration management tasks, you can also view the BPEL Processes, its different versions, and the suitcase files associated with each version. You can also compare the BPEL Process suitcase files of different versions and track the changes that were made to a version. This allows you to track changes and identify the cause for changed behavior of the BPEL Process.</li> </ul>	<a href="#">Creating Infrastructure and Aggregate Services</a> <a href="#">Adding SOAP Tests for Partner Links</a>
	<p>Business Process Execution Language (BPEL) is an XML-based language for enabling task sharing across multiple enterprises using a combination of Web services. Oracle BPEL Process Manager provides a framework for easily designing, deploying, monitoring, and administering processes based on BPEL standards.</p>	

**Table 8–1 (Cont.) Middleware Targets Monitored**

Monitored Target	Services Offered by Grid Control	Related Links
Oracle Application Server Farm	<ul style="list-style-type: none"> <li>■ Helps you to manually discover Oracle Application Server Farms and monitor their overall health.</li> <li>■ Helps you to view the member DCM clusters and application servers within a farm.</li> <li>■ Also helps you to view the topology and perform administrative tasks like creating jobs and blackouts, and searching for deployments and configuration information.</li> </ul> <p>An OracleAS Farm is a collection of OracleAS DCM Managed Clusters and Oracle Application Server instances that share the same Farm Repository.</p>	<a href="#">Viewing Members</a> <a href="#">Viewing Topology</a> <a href="#">Monitoring Dashboard</a>
Oracle Application Server DCM Managed Cluster and Oracle Application Server Clusters	<ul style="list-style-type: none"> <li>■ Helps you to manually discover Oracle Application Server DCM Managed Clusters and Oracle Application Server Clusters, and monitor their overall health.</li> <li>■ Helps you to study the High Availability grouping done for the members of an DCM Managed Cluster. A High Availability Group is a group composed of similar individual components of application server instances clustered together in a DCM managed cluster. For example, an OC4J High Availability Group has a group of OC4J instances in an OracleAS Cluster.</li> <li>■ Also helps you to view the topology and access the Oracle System Monitoring Dashboard that provides a quick overview of the health of the target.</li> </ul> <p>An OracleAS DCM Managed Cluster is a collection of release 9.0.4 to release 10.1.2 Oracle Application Server instances with identical configuration and application deployment characteristics. An OracleAS Cluster is a collection of release 10.1.3 or higher versions of Oracle Application Server instances.</p>	<a href="#">Viewing Members</a> <a href="#">Viewing Topology</a> <a href="#">Monitoring Dashboard</a>



**Table 8–1 (Cont.) Middleware Targets Monitored**

Monitored Target	Services Offered by Grid Control	Related Links
Oracle BI Suite Enterprise Edition (Oracle BI Suite EE)	<ul style="list-style-type: none"> <li>■ Helps you to manually discover Oracle BI Suite EE targets, and monitor their overall health. Also help you to manually discover Data warehouse Application Console Server (DAC Server).</li> <li>■ Allows you to diagnose, notify, and correct performance and availability problems in Oracle BI Suite EE targets.</li> <li>■ Allows you to access current and historical performance information using graphs and reports. This help you to perform a root cause analysis to drill down to the problem areas and fix them before they affect the end users.</li> <li>■ Also helps you to perform enterprise configuration management tasks like viewing, comparing, and searching configuration information.</li> </ul> <p>Oracle BI Suite EE is an integrated suite of products that offers a comprehensive and integrated set of business intelligence tools that help you collect information from a variety of sources, analyze it, and share it with the broadest audiences of users. It is a logical, composite target and its components are Oracle BI Presentation Server, Oracle BI Cluster Controller, Oracle BI Analytics Server, and Oracle BI Scheduler. Additionally, you can monitor DAC Server.</p>	
Oracle Service Bus	<ul style="list-style-type: none"> <li>■ Automatically discovers and monitors Oracle Service Bus targets deployed to Oracle WebLogic Managed Servers.</li> <li>■ Allows you to create infrastructure services for Oracle Service Bus targets, and SOAP Tests for the proxy services.</li> <li>■ Allows you to run deployment procedures for provisioning Oracle Service Bus resources.</li> </ul>	<p><a href="#">Creating Infrastructure and Aggregate Services</a></p> <p><a href="#">Adding SOAP Tests for Oracle Service Bus Proxy Services</a></p>

**Table 8–1 (Cont.) Middleware Targets Monitored**

Monitored Target	Services Offered by Grid Control	Related Links
Third-Party Application Server	<p>Helps you manually discover the following third-party application servers, and monitor their status:</p> <ul style="list-style-type: none"> <li>■ IBM WebSphere Application Server Cell</li> <li>■ IBM WebSphere Application Server Cluster</li> <li>■ IBM WebSphere Application Server</li> <li>■ JBoss Application Server</li> <li>■ Microsoft Exchange Organization</li> <li>■ Microsoft Exchange Routing Group</li> <li>■ Microsoft Exchange Server</li> </ul> <p>Also helps you view the applications deployed on these third-party application servers, and perform enterprise configuration management tasks like viewing, comparing, and searching configuration information. Helps you view the members within a domain, cluster, or cell.</p>	Chapter 15: Host and Third-Party Target Management
Third-Party Plug-In	<p>Helps you manually discover the following third-party plug-ins and monitor their health:</p> <ul style="list-style-type: none"> <li>■ IBM WebSphere MQ Cluster</li> <li>■ IBM WebSphere MQ Queue Manager (along with the associated Queues and Channels)</li> </ul>	Chapter 15: Host and Third-Party Target Management

---

**Note:** For more information about monitoring third-party targets, see **Chapter 15: Host and Third-Party Target Management**.

---



---

**Important:** Enterprise Manager 10.2 release supported only release 9.0.4 to release 10.1.2 application server targets, and for that release, OracleAS DCM Managed Cluster was actually called "OracleAS Cluster". However, now that Enterprise Manager 10.2.0.2 release has additional support for 10.1.3 application server targets, for the purpose of clarity, the terminologies have been amended to distinguish the ones managed by DCM and the ones that relate to release 10.1.3 application server targets.

---

## Viewing Critical Information

This section describes how you can use Grid Control to view critical information pertaining to middleware targets.

### General Information

Grid Control provides general information about the targets including their status and availability. This helps you understand how the target is performing, where it is deployed, what is its version, where is its home directory, and so on. If the target is a logical target that groups its inherent members, then details about the membership are also provided.

**Figure 8–2 General Section**

**General**

 Status **Up** Black Out

Up Since **Oct 25, 2006 3:23:43 PM**

Availability (%) **100**  
(Last 24 Hours)

Oracle Home [/private/home/oracle/102gc/oms10g](#) 


Host [stang13.us.oracle.com](#)

## Alerts and Host Alerts



Grid Control shows the number of critical, warning, and error alerts generated for the past 24 hours. These alerts indicate that a particular metric condition has been encountered. For example, an alert is triggered when a metric threshold is reached. Using these details, you can drill down to investigate the target and the problem that triggered the alert. For more information about alerts, see [Automated Monitoring and Alerts](#) in this chapter.

**Figure 8–3 Alerts and Host Alerts Sections**

### Alerts

Metric	Severity	Alert Triggered	Last Value	Last Checked
<a href="#">Error Rate (%)</a>		Nov 8, 2006 7:25:08 PM	3.26	Nov 8, 2006 7:25

### Host Alerts

Metric	Severity	Alert Triggered
<a href="#">Filesystem Space Available (%) for /</a>		Oct 4, 2006 12:58:06 AM
<a href="#">Log File Pattern Matched Line Count for /private/oracle/orahomes/agent10g/sysman/log/emagent.log/ERROR/Sun Oct 22 16:37:15 2006 GM -07:00</a>		Oct 22, 2006 4:37:14 PM

## Policy Violations

Grid Control also shows the number of informational, warning, and critical policy rules violated for the application server target. You can see the roll up of individual policy compliance scores for that target. Compliance scores provide a quick way to determine the health of the your middleware target. You can also see the policy trend overview for the last 24 hours, for the last week, for the last month, or for a user-define time period, and determine your course of action in solving the policy violations.

Grid Control also shows the time when security policy rules were last evaluated, and shows the roll up of individual policy compliance scores for that target.

**Figure 8–4 Policy Violations and Security Sections**

### Policy Violations

Current **1**    Distinct Rules Violated **1**    [Compliance Score \(%\) 17](#) [Policy Trend Overview](#)

### Security

Last Security Evaluation  **Nov 8, 2006 7:35:23 PM PST** [Compliance Score \(%\) 17](#) [Enterprise Se...](#)

## Viewing Information About Applications Deployed

Grid Control provides details about the J2EE applications deployed on the application server instances. You can quickly assess which applications are most active and which are using the most system resources.

**Figure 8–5 Viewing Details about the Applications Deployed**

Oracle WebLogic Managed Server: Server-10

Data Retrieved Jan 7, 2009 4:23:42 AM PST View Data Last 24 hours

Home Applications Web Services Resources Performance Administration

Create Service Start Stop

Select	Application Name	Status	Average Active Sessions	Average Requests Per Second	Average Response Time
<input checked="" type="radio"/>	<a href="#">bea-guardian-agent</a>		0.00	0.00	
<input type="radio"/>	<a href="#">lodgingsupplier</a>		0.00	0.00	
<input type="radio"/>	<a href="#">activitiesupplier</a>		0.00	0.00	
<input type="radio"/>	<a href="#">opc</a>		0.00	0.00	
<input type="radio"/>	<a href="#">airlinesupplier</a>		0.00	0.00	
<input type="radio"/>	<a href="#">bank</a>		0.00	0.00	
<input type="radio"/>	<a href="#">consumerwebsite</a>		0.00	0.00	

Home Applications Web Services Resources Performance Administration

Grid Control gives you the flexibility to update the information about applications based on current, real-time application performance, or based on historical data from the Management Repository. For example, you can see which applications were the most active over the past 24 hours, over the past week, or over the past month.

## Viewing Members

For logical, composite targets such as Oracle WebLogic Server Domains, OracleAS Farms and Clusters, and so on, Grid Control provides rolled-up information on the member targets that make up the group. It provides a "Members" tab that helps you view the members associated with that particular composite target; check their status; and start, stop, or restart each of them.

**Figure 8–6 Oracle WebLogic Server Domain**

Oracle WebLogic Server Domain: medrec\_7011

Page Refreshed Jan 7, 2009 4:19:45 AM PST Refresh Launch Dashbo

Home Administration Members

Search All  Go View Flat

Start Stop Restart

Expand All Collapse All

Select	Name	Type	Status	Status Details	Alerts	Vis
<input checked="" type="radio"/>	<a href="#">stbdm01.us.oracle.com.medrec_7011</a>	Oracle WebLogic Server Domain	n/a	3 (↑3)		
<input type="radio"/>	<a href="#">stbdm01.us.oracle.com.medrec_7011.myCluster</a>	Oracle WebLogic Server Cluster		2 (↑2)		
<input type="radio"/>	<a href="#">stbdm01.us.oracle.com.medrec_7011.myCluster.Server-10</a>	Oracle WebLogic Managed Server		↑		
<input type="radio"/>	<a href="#">stbdm01.us.oracle.com.medrec_7011.myCluster.Server-9</a>	Oracle WebLogic Managed Server		↑		
<input type="radio"/>	<a href="#">stbdm01.us.oracle.com.medrec_7011.MedRecServer</a>	Oracle WebLogic Managed Server		↑		

Grid Control also helps you study the High Availability (HA) grouping done for the members of an OracleAS DCM Managed Cluster. A High Availability Group is a group composed of similar individual components of application server instances clustered together in a DCM managed cluster. For example, an OC4J High Availability Group has a group of OC4J instances in an OracleAS Cluster.

## Monitoring Performance

This section describes the features offered by Grid Control that help you monitor the performance of middleware targets.

### Using Performance Metrics

Grid Control provides a set of predefined performance metrics for every middleware target. By selecting a metric, you can determine whether the thresholds have been defined for a particular metric. Grid Control uses thresholds defined within metrics as a mechanism to generate alerts. These alerts in turn are used to notify you whether a target is up or down, a disk on the target is near full, and so on. Thus, you can monitor their overall performance.

The performance metrics provide details about the metric as a current real time value (30 seconds, 1 minute, or 5 minutes) or a previous value (past 24 hours, 7 days, or 31 days). The historical information is displayed as graphs and a table. By using graphs, you can easily watch for trends, and by using tables, you can examine details of past metric severity history.

### Automated Monitoring and Alerts

Grid Control provides a comprehensive set of features that facilitates automated monitoring and generation of alerts. Oracle Management Agent on a host automatically discovers the middleware targets on that host, and helps Grid Control perform unattended monitoring of their status, health, and performance.

Grid Control gathers and evaluates diagnostic information from these targets distributed across the enterprise, and an extensive array of middleware performance metrics are automatically monitored against predefined thresholds.

For example, Grid Control can automatically monitor:

- The CPU or memory consumption of the application server, including detailed monitoring of individual Java Virtual Machines (JVMs) being run by the server's Oracle Application Server Containers for J2EE (OC4J) instances
- J2EE application responsiveness from the application down through individual servlets and Enterprise JavaBeans (EJBs)
- HTTP Server session volumes, connection duration, and error rates
- Oracle Application Server Web Cache hit rates and volumes
- Top servlets based on number of requests, maximum processing time, and highest average processing time

If an Oracle Application Server or any of its core components go down, or if a performance metric crosses a warning or critical threshold, then an alert is generated by Grid Control and a notification is sent to you. Grid Control supports notifications via e-mail (including e-mail-to-page systems), SNMP traps, and/or by running custom scripts.

When you receive an alert notification, Grid Control makes it easy for you to investigate the problem and take corrective actions wherever required. For example, notification of excessive CPU consumption by OC4J may lead to investigation of the applications running in that container. By using the Top J2EE Applications tab of the Application Server Home page (Figure 8-7) in Grid Control, you can quickly identify the highest volume or least responsive application. You can then drill down and diagnose application's servlets, Java Server Pages (JSPs), or EJBs to identify the bottleneck.

**Figure 8–7 Monitoring the Performance of Your Deployed J2EE Applications**

Select	Name	OC4J Instance	Total Processing Time (seconds)	Average Servlet/JSP Processing Time (seconds)	Servlet/JSP Requests Processed	T Servlet/Process T (secor)
<input checked="" type="radio"/>	system	home	10.64	0.01	1,102	10
<input type="radio"/>	policymanager	home	7.93	0.03	264	7
<input type="radio"/>	default	home	1.03	0.002	477	1
<input type="radio"/>	gateway	home	0.00	0.00	0	0
<input type="radio"/>	coreman	home	0.00	0.00	0	0
<input type="radio"/>	SelectManufacturingUI	home	0.00	0.00	0	0

You can set up corrective actions to automatically resolve an alert condition. These corrective actions ensure that routine responses to alerts are automatically executed, thereby saving you time and ensuring that problems are dealt with before they noticeably impact the users.

You can also use monitoring templates to simplify the task of standardizing monitoring settings across your enterprise. You can specify the monitoring settings once and apply them to all Oracle Application Server targets. A Monitoring template defines all Grid Control parameters you would normally set to monitor an Oracle Application Server target, such as:

- Target type to which the template applies
- Metrics (including user-defined metrics), thresholds, metric collection schedules, and corrective actions

When a change is made to a template, you can reapply the template across affected Oracle Application Server targets in order to propagate the new changes. You can reapply monitoring templates as often as needed.

## Viewing Topology

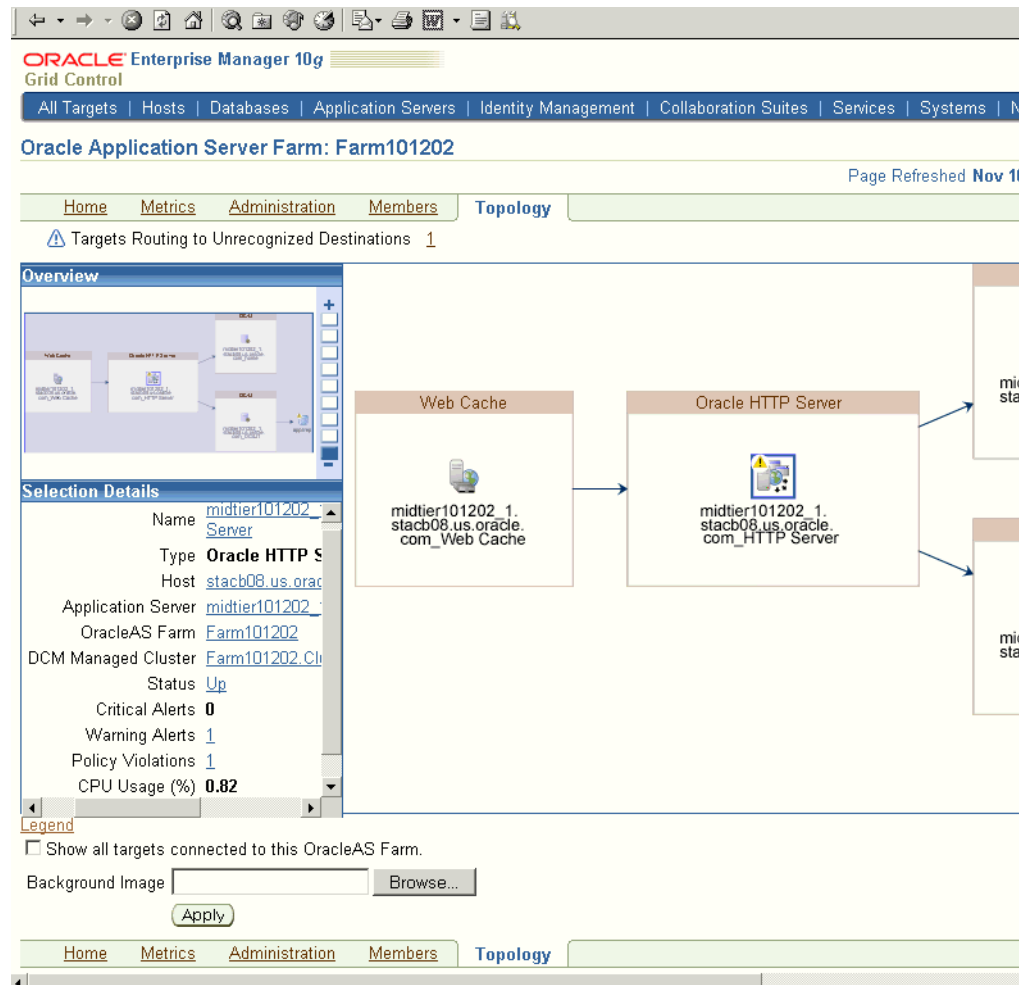
Besides the monitoring features provided for Oracle Application Servers, OracleAS Farms, and OracleAS DCM Managed Clusters, Grid Control also provides a topology view for these targets. You can view their topology to understand what application servers and components are running on which hosts, how these components are related to each other, and how requests are routed through different layers of the deployment. This kind of visualization of the data center enterprise topology helps administrators effectively monitor, manage, and validate the enterprise architecture.

Grid Control provides three different views of topology. Each view provides the overlay of some key metrics of components including current status, number of alerts and policy violations, and CPU/memory utilization performance metrics.

- **Host View** shows the physical view of deployment and visually shows the relationship between hosts and various components and instances hosted by them.
- **Routing Overview** shows the routing view of topology and provides an end-to-end view of wired component and application flow through them. The routing includes routing from OracleAS Web Cache to Oracle HTTP Server to OC4J to DB instances.

- **Routing Details** also shows the protocol and port used by various components to route requests to other components in the topology.

**Figure 8–8 Topology View of an OracleAS Farm**



These routing views enable you to fix various topology configuration problems immediately. For example, if an OC4J instance is not accepting any user requests, the routing details view confirms if the front-ending Oracle HTTP Server instances are configured correctly to route user requests to that OC4J instance. OracleAS Web Caches, Oracle HTTP Server, or OC4J components are displayed in the same box, if they provide redundancy for each other.

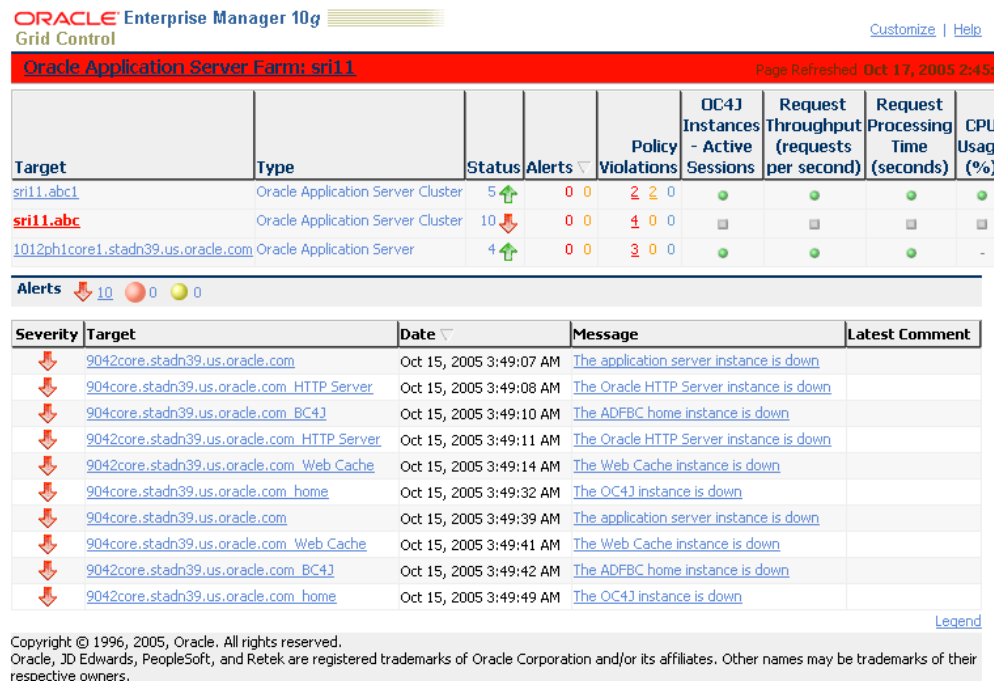
For example, a set of OracleAS Web Caches are combined in the same box if all of them are routing to the same set of Oracle HTTP Server instances. Similarly, a set of OC4J instances are combined in the same box if they are routed by the same set of Oracle HTTP Server instances, and are also hosting the same set of J2EE applications. As the like-components are grouped together, the topology visual representation in routing overview or routing details view provides instant Root Cause Analysis for the service availability problems.

## Monitoring Dashboard

Grid Control also provides access to Oracle System Monitoring Dashboard that helps you monitor the health of a target. However, this feature is provided only for OracleAS Farms, OracleAS DCM Managed Clusters, Oracle Application Servers, OC4J High Availability Groups, and HTTP Server High Availability Groups.

Oracle System Monitoring Dashboard presents information using intuitive icons and graphics that let you spot recent changes and quickly identify and respond to problems. You can customize the display attributes to match information requirements of managed targets, monitor status indicators for recent problems, and see new alerts that have been triggered since the dashboard was last viewed.

**Figure 8–9 Oracle System Monitoring Dashboard**



## Administering Middleware Targets

Grid Control provides features for performing administrative tasks on middleware targets. It provides Web-based interfaces for performing operations such as:

- Starting, stopping, or restarting middleware targets.
- Performing configuration management tasks, such as viewing and comparing configuration information. Refer to "[Managing Configurations](#)" in this chapter for more information.
- Accessing a list of predefined search queries to search and retrieve configuration information. Refer to "[Managing Configurations](#)" in this chapter for more information.
- Integrating application instrumentation in Enterprise Manager's event monitoring infrastructure. Refer to "[Extensible Monitoring](#)" in this chapter for more information.



- Staging or applying an interim patch or patch set, and/or cloning an application server's Oracle home to one or more hosts. Refer to "[Cloning and Patching the Middleware Environment](#)" in this chapter for more information.
- Performing scheduled backup and recovery. Refer to "[Managing Backup and Recovery of the Application Server Environment](#)" in this chapter for more information.
- Creating blackouts to perform scheduled maintenance.
- Viewing the number of scheduled, running, suspended, and problem (stopped/failed) executions for all Grid Control jobs submitted on the application server.
- Creating Application Server Farms and managing their components.

## Managing and Creating Blackouts

Grid Control comes with a bundle of performance and health metrics that enable automated monitoring of application targets in your environment. When a metric reaches the predefined warning or critical threshold, Grid Control generates an alert and notifies the administrators.

However, there are occasions when you want to perform maintenance work on your middleware targets, and not want any alerts to be generated while you are bringing them down. In this case, you can schedule a blackout and suspend monitoring of the middleware targets.

Blackouts allow you to suspend any data collection activity on one or more monitored targets, thus allowing you to perform scheduled maintenance on targets. If you continue monitoring during these periods, the collected data will show trends and other monitoring information that are not the result of normal day-to-day operations. To get a more accurate, long-term picture of a target's performance, you can use blackouts to exclude these special-case situations from data analysis.

Grid Control allows you to define new blackouts; view the status of existing blackouts; and edit, stop, and delete blackouts that are not required. You will find the Black Out option in the General section on the home page of any application server target. For composite targets like OracleAS Farms, you will find this option in the Administration tab.

## Managing and Creating Jobs

For OracleAS Farms, OracleAS DCM Managed Clusters, OracleAS Clusters, Oracle WebLogic Server Domains, and Oracle WebLogic Server Clusters, Grid Control provides a job system that allows you to create jobs. For example, for an Oracle WebLogic Server Domain, you can create a job to automatically start, stop, or restart that server. You can also view details about the jobs that are scheduled, running, suspended, or the ones that have a problem. You can schedule a job directly from these composite targets to run an OS command on the members or start, stop, or restart any OPMN component.

## Creating Services from J2EE or Web Applications

Grid Control allows you to create services directly from a J2EE or Web application to help you track their availability and performance. Understandably, the business applications running on an application server simply reflect the quality of service rendered by an organization. Considering their criticality and complexity, it becomes

imperative for any organization to ensure that they are always available and capable of servicing the requests within an acceptable turnaround time.

By creating services from a J2EE or Web application, you can group your applications as an entity that provides a useful function to its users. This helps you monitor the overall performance of the service that represents a function offered to the users, and also monitor the infrastructure components upon which the service depends. You can also receive alerts when there is a problem, identify common issues within the system, diagnose causes of failures, and resolve them.

**Figure 8–10 Creating Service from J2EE or Web Applications**

Select Name	OC4J Instance	Total Processing Time
<input checked="" type="radio"/> system	home	
<input type="radio"/> policymanager	home	
<input type="radio"/> default	home	

Every time you create a service for J2EE or Web application, Grid Control creates a service dependency to identify the components on which the selected application depends. Here, the components refer to infrastructure components like hosts, databases, application servers, and so on that work together to host your J2EE or Web applications. After identifying the components required for that application, Grid Control creates a system with those components, and then creates a service and associates it with that application.

So, every time you create a service, Grid Control checks to see if a dependency mapping is already available for the selected application. If it is available, then a service is created using that existing system rather than creating another system with exactly the same members.

According to the availability of this dependency mapping for an application, you can either create a new service with a new system, create a new service with an existing system, or refresh a service.

## Creating Infrastructure and Aggregate Services

Grid Control allows you to create infrastructure services for middleware targets such as Oracle BPEL Process Manager targets and Oracle Service Bus targets.

---

**Note:** "Middleware targets" in this section refer to Oracle BPEL Process Manager targets and Oracle Service Bus targets only.

---

An infrastructure service is a dependency service that is created to identify the infrastructure components on which the middleware target depends. Here, the infrastructure components refer to hosts, databases, application servers, and so on that work together to host the middleware target.

You can either create an infrastructure service with a new system or an existing system, or simply refresh an existing infrastructure service, if there is already one

existing. By creating infrastructure services and systems, you can better manage your middleware targets and also the components on which the middleware targets depend.

**Figure 8–11 Creating Infrastructure Service for Oracle BPEL Process Manager**

**System And Service**

Refresh BPEL Infrastructure Service

System	Status	Alerts	Policy Violations	Service
SOA_Instance_system	↓	6 0	5 2 0	SOA_Instance

For example, once you create an infrastructure service for an Oracle BPEL Process Manager target, Grid Control allows you to create an aggregate service for every process within that Oracle BPEL target. An aggregate service is a logical grouping of services, in this case, infrastructure services and availability services. An availability service is a service that is created when a SOAP (Simple Object Access Protocol) test is added for the first time to the partner links associated with the selected Oracle BPEL processes. Aggregate Services give you a bird's-eye view of the services that have been created for the Oracle BPEL target and helps you monitor their availability, performance, and usage.

**Figure 8–12 Creating Aggregate Service for Oracle BPEL Process**

**Services**

Provides details about the services that have been created for this Oracle BPEL Process.

Create Service

Expand All | Collapse All

Name	Service Type	Status	Usage Alert	Performance Alert	Policy Violations
▼ default_BatchOrderProcessing(v.1.0)	Aggregate Service	↓	0 0	0 0	0 0 0
▶ SOA_Instance	Generic Service	↓	0 0	0 0	0 0 0

## Adding SOAP Tests for Partner Links

Grid Control allows you to add SOAP tests for partner links associated with a selected Oracle BPEL process in order to test their availability.

SOAP is an XML-based message protocol used by Web Services of these partner links. Each SOAP message consists of a single SOAP envelope. An envelope defines how to process the message, who should process the message, and whether processing is optional or mandatory.

Grid Control allows you to select a particular method specific to a port type exposed by the Web service, and provide values for the input parameters that it should use. The SOAP test invokes the selected method and checks if the Web service is up. This way, you can track the performance of the partner links and know when they are down.

## Adding SOAP Tests for Oracle Service Bus Proxy Services

Grid Control allows you to create a SOAP test for either a proxy service or an end point associated with the business service of an Oracle Service Bus. The SOAP tests that are added help in monitoring the availability and response time of the Web service used by these partner business services.

When a SOAP test is added for the first time, an availability service for it is created automatically. This availability service gets added as a child service of the aggregate service. Every other SOAP test that is added gets included in the same availability service.

## Adding SOAP Tests for Web Services of Applications Deployed to Oracle WebLogic Managed Server

Grid Control allows you to create SOAP tests for web services of applications deployed to Oracle WebLogic Manager Servers. When a SOAP test is added for the first time, a generic service for it is created automatically. The generic service is created to monitor the availability of the web services in the application. By default, the generic service that is created has the format <appname>\_webservices\_ availability. The next time you create a SOAP test, Grid Control associates the new SOAP test with the existing service.

## Provisioning BPEL Processes

Grid Control provides deployment procedures for provisioning BPEL processes on selected BPEL Process Managers. The BPEL Process Provisioning deployment procedure comes with the following predefined phases:

- **Stage Setup:** This phase sets up the BPEL processes and other dependencies. Essentially, it sets up the BPEL suitcase files, stages them to a temporary location for deployment, and stages the deploy.jar on the target hosts.
- **Deploy:** This phase deploys the BPEL processes on the selected target BPEL Process Managers.
- **Clean Setup:** This phase cleans up the files and folders staged on the individual hosts.

The BPEL process suitcase files have to be placed as generic components in Oracle Software Library, and once it is available for deployment, the deployment procedure can be run to deploy the processes on BPEL Process Managers.

## Provisioning Oracle Service Bus Resources

Grid Control provides deployment procedures for deploying exported Oracle Service Bus resources to a target Oracle Service Bus domain, and move your configuration from one environment to another.

Oracle Service Bus resources can be organized into individual projects. Projects are non-hierarchical, disjointed, top-level grouping constructs. All resources (such as business services, proxy services, WS-Policies, WSDLs, schemas, XQuery transformations, JARs, and so on) reside in exactly one non-overlapping project. Resources can be created directly under a project or be further organized into folders. Folders may be created inside projects or inside other folders, and the folders are similar to directories in a file system, with the project level being the root directory.

Using the *Oracle Service Bus Resource Provisioning* deployment procedure, you can select the projects of an Oracle Service Bus domain that you want to export, and deploy the resources of that project to another Oracle Service Bus domain. You can export the resources at a project level or resource level. Either way, technically, the resources are exported as a JAR file, and the JAR file is deployed to the target Oracle Service Bus domain. If you already have the resources exported as a JAR file from an Oracle Service Bus domain, then you can upload the file to Oracle Software Library

(software library) and use the deployment procedure to deploy it from the software library.

The Oracle Service Bus Resource Provisioning deployment procedure offers the following predefined phases:

- **Export Resources:** Exports all the resources from the source Oracle Service Bus domain.
- **Save Exported JAR File:** Saves the exported JAR file as a generic component in Oracle Software Library (Software Library).
- **Transfer Exported JAR File for Deployment:** Transfers the JAR file from source host to the destination for deployment.
- **Stage Selected Software Library Components for Deployment:** Copies the selected Software Library components to the host of the Oracle Service Bus domain where they have to be deployed.
- **Deploy Resources and Customize:** Deploys resources and customizes on the selected target.

You can also use a customization file to modify the environment values and references within resources to suit the changed environment. Customization files can include customizations for all the environment values found in the selected resources. In addition, it can also include changes to the reference customization type for changing references inside resources with dependencies. You can create these customization files using the Oracle Service Bus console.

---

---

**Note:** Oracle Service Bus provisioning is not supported if Oracle WebLogic Server Domain is discovered using a remote Management Agent. For example, if you use the Management Agent running on host2 to discover Oracle WebLogic Server Domain and Oracle Service Bus running on host1, and then if you try provisioning Oracle Service Bus resources, you will see errors.

---

---

## Creating Log Rotation Jobs

Oracle Application Server components generate log files containing messages that record all types of events, including startup and shutdown information, errors, warning messages, access information on HTTP requests, and additional information.

However, the information recorded in log files is voluminous, thus making it difficult to track what update was made at what time. Also because of the huge quantity of information updated periodically, the log files grow in size and occupy more space on the system over a period of time. The only way to manage these log files is to manually archive the contents to another file and store them in a different location.

Considering these impediments, Grid Control has been enhanced with a log rotation feature that helps you manage the logs of Oracle Application Server components more effectively. In particular, you can use Grid Control to:

- Schedule a job that automatically rotates a log at the scheduled date and time
- Manage space on your system by storing the rotated log files in a different directory

Grid Control allows you to view the logs of a particular Oracle Application Server component type and select the ones that need to be rotated. Note that a log rotation job can also be part of a multi-task job.

When a log rotation job is executed, Grid Control automatically stops the component whose logs have to be rotated. After it is stopped, the content from its existing log file is moved to another file that is distinguished with the timestamp when it was actually rotated. The original log file is kept empty for new log details to be populated. Once this is done, Grid Control restarts the component.

## Managing Configurations

Grid Control provides a suite of configuration management functions that can be performed on middleware targets.

Oracle Management Agent collects configuration information about Oracle Application Server targets from their respective configuration files, and communicates this information over HTTP/HTTPS to Oracle Management Service, which stores it in the Management Repository. This information is periodically collected and updated while maintaining the audit of changes. Enterprise Manager's configuration management capabilities efficiently guide the users to desired configuration data in a particular component.

**See Also:** "Hardware and Software Configurations" in [Chapter 4, "Enterprise Configuration Management"](#)

You can compare these configuration details and view the differences and similarities between the two instances of a middleware target. You have the flexibility to compare two last collected configurations or two saved configuration files. You can also compare one configuration with multiple configurations or one configuration in the Management Repository with a saved configuration file.

Using Grid Control, you can search configurations across middleware targets and find configuration anomalies - whether they are a mismatch of an install/patch version of Oracle Application Server software, or they are a mismatch of software configuration data for the core components of Oracle Application Server. You can perform more intelligent searches to identify all the components hosting a particular application or other resources.

For example, the Administration page in Grid Control provided for Oracle Application Server targets allows you to search one or more:

- Origin servers
- Application servers with particular installation settings
- Data sources used by the applications deployed across your enterprise configuration
- J2EE applications deployed in a particular OC4J instance, application server instance, or host
- Modules of J2EE applications deployed across your enterprise configuration
- Application server ports across your enterprise topology

In addition, for BPEL Process Manager targets, you can view the BPEL Processes, its different versions, and the suitcase files associated with each version. You can also compare the BPEL Process suitcase files of different versions and track the changes that were made to a version. shows you how the versions can be selected and

compared. This allows you to identify the cause for improved or deteriorated performance due to a change in the BPEL Process suitcase file.

**Figure 8–13 BPEL Process Suitcase File Comparison**

**Compare Process Versions: HelloWorld**

First Collected	<a href="#">as13.iwinrec22.us.oracle.com_bpel</a> Sep 6, 2007 10:10:54 PM PDT	Second Collected	<a href="#">as13.iwinrec22.us.oracle.com_bpel</a> Sep 6, 2007 10:10:54 PM PDT
Version	1.0	Version	1.0

**TIP** The comparison results below show the difference for the selected process version. Use the Process Version menu to compare file contents, select a file and click Compare File Contents.

Compare File Contents		Sep 6, 2007 10:10:54 PM PDT	
Select	Result	Filename	Filename
<input checked="" type="radio"/>		<a href="#">HelloWorld.wsdl</a>	<a href="#">HelloWorld.wsdl</a>
<input type="radio"/>		<a href="#">bpel.xml</a>	<a href="#">bpel.xml</a>
<input type="radio"/>		<a href="#">HelloWorld.bpel</a>	<a href="#">HelloWorld.bpel</a>

## Extensible Monitoring

Many administrators often require custom logic to be written to check for conditions specific to their application environments. Grid Control allows integration of application instrumentation in Grid Control's event monitoring infrastructure. If application developers expose application instrumentation using standards like JMX or Web Services operations, then you can build management plug-ins for the instrumentation using easy-to-use command line tools, and leverage Grid Control's event monitoring system to monitor it. You do not have to edit any XM files or write any integration code to integrate such instrumentation in Grid Control. Follow these procedures to integrate application-defined instrumentation in Grid Control:

- Use Command Line Interfaces that analyze MBean interfaces for JMX and WSDL for Web Services and create management plug-ins
- Import Management Plug-in Archive in Grid Control
- Deploy Management Plug-in to Management Agents
- Create Target-type instances for the target types defined in Management Plug-in Archive
- Leverage Grid Control's event monitoring system including monitoring templates, corrective actions, historical and real time metric views, alerts, customization of notification rules, and methods on events generated from application instrumentation metrics

## Diagnosing Performance Issues with Top Reports

When you are troubleshooting performance problems, it can be helpful to know which servlets or JSPs are the most active. By using the Top Servlets or Top JSPs performance links of the Application Server Performance page (Figure 8–14 and Figure 8–15) in Grid Control, you can identify the top Java servlets or JSPs running on the application server instance. You can then sort them to identify the servlets and JSPs by number of requests, maximum processing time, or highest average processing time.

**Figure 8–14 Monitoring the Top Servlets for a J2EE Application**

Top Servlets

Page Refreshed Aug 17, 2005 11:31:50 PM PDT Refresh View Data Real Time: Manual Refresh

View Top Servlets by Requests Processed

Name	OC4J Instance	Application	Web Module	Average Processing Time (seconds)	Requests Processed	Total Processing Time (seconds)
upload	OC4J_EM	em	em	0.05	10,606	509.79
Spy	OC4J_EM	default	dms	0.002	9,130	21.92
OHW	OC4J_EM	em	em	0.29	699	202.60
jsp	OC4J_EM	em	em	0.22	534	117.27
console	OC4J_EM	em	em	0.75	488	364.86
dynamicImage	OC4J_EM	em	em	0.12	210	25.85
logon	OC4J_EM	em	em	0.18	61	11.06
TopoServlet	OC4J_EM	em	em	0.006	2	0.01
SVGServlet	OC4J_EM	em	em	22.83	1	22.83
genwallet	OC4J_EM	em	em	0.008	1	0.008

**Figure 8–15 Monitoring the Top JSPs for a J2EE Application**

Top JSPs

Page Refreshed Aug 17, 2005 11:52:59 PM PDT Refresh View Data Real Time: Manual Refresh

View Top JSPs by Requests Processed

Name	OC4J Instance	Application	Web Module	Average Processing Time (seconds)	Requests Processed	Total Processing Time (seconds)
ip/render/elem/tableRender.jsp	OC4J_EM	em	em	0.001	38	0.05
ias/faintTabsInclude.jsp	OC4J_EM	em	em	0.0003	29	0.01
ias/faintTabsIncludeBottom.jsp	OC4J_EM	em	em	0.0004	28	0.01
home.jsp	OC4J_EM	em	em	0.41	24	9.86
ip/render/justInTimeSelectTargets.jsp	OC4J_EM	em	em	0.04	22	0.81
ip/render/previewReport.jsp	OC4J_EM	em	em	0.36	19	6.92
ip/render/elem/chartRender.jsp	OC4J_EM	em	em	0.003	19	0.05
ip/render/elem/textRender.jsp	OC4J_EM	em	em	0.0009	16	0.01
target/selectTargetSelection.jsp	OC4J_EM	em	em	0.06	15	0.97
ip/reportDefinitionList.jsp	OC4J_EM	em	em	0.004	14	0.06

## Analyzing Historical Performance

As with all Grid Control diagnostics, the application server diagnostic reports can be based on current or historical data. Application server metrics are collected and stored in the Management Repository, so you can analyze the data well after the situation has changed. For example, you can use historical data and diagnostic reports to research an application performance problem that occurred days or even weeks ago.

You can even provide a customized time period for which the data should be retrieved from the Management Repository. You can customize the time period for:

- Pre-defined range of the last 24 hours, last 7 days, or last 31 days
- Customized range of any number of days, weeks, months, or years
- Any start date and end date (such that the duration is not greater than 99 years)

## Maintaining Middleware Targets

This section describes the features offered by Grid Control that help you maintain middleware targets.

## Managing Backup and Recovery of the Application Server Environment

Backup and recovery refers to the various strategies and procedures involved in guarding against hardware failures and data loss, and reconstructing data should there be a loss. A comprehensive backup strategy should involve a coordinated approach to backing up your entire application server environment, including the middle tiers and the Application Server Infrastructure Oracle homes.

Grid Control helps you manage the backup and recovery of a single application server or a group of application servers.

Using Grid Control, you can:



- Schedule backups
- Restore application server backups for recovery
- Display the status of backup jobs
- Display the status of recovery jobs
- Configure the required settings for backup

For more information on the types of backups and recommended strategy for performing backup and recovery, see the *Oracle Application Server Administrator's Guide*.

## Cloning and Patching the Middleware Environment

Using Grid Control's automated provisioning tools, you can ensure standardization in your data center and also significantly reduce the time spent on these tasks. To consistently maintain standardization in the topology, it is recommended that the new instances be added through "cloning" rather than "install and configure".

Cloning ensures that the new instance is installed and configured exactly like other instances in the enterprise topology. Cloning is the process of copying an existing installation to a different location while preserving its configuration and deployments. Grid Control's cloning wizard automates the duplication of application server [About Access and Identity Management](#) installations; specifically, the directories where the Oracle homes reside. Its "multicasting" capability also helps you create multiple clones on multiple target hosts in a single operation.

Using a direct link to My Oracle Support, Grid Control proactively and regularly retrieves the list of critical patches that have to be applied on Oracle Application Server installations. Grid Control also analyzes the data center environment and notifies you of patches that are applicable to their application server instances. All other patches can also be manually found in the context of a specific target. You can also automate the application of patches using robust job system infrastructure. You can apply the patches instantly or schedule the application in the maintenance window while backing out Oracle Application Server instances in the maintenance window.

**See Also:** "Cloning" and "Patching" in [Chapter 13, "Lifecycle Management"](#)

## Managing Identity Management Targets

This section describes how you can use Grid Control to manage your Identity Management targets. It describes the following:

- [About Access and Identity Management](#)
- [Using Grid Control for Monitoring Identity Management Targets](#)
- [Identity Management Systems](#)
- [Automated Identity Management and Alerts](#)
- [Identity Management Root Cause Analysis](#)
- [Diagnosing Identity Management Performance and Availability Problems](#)
- [Leveraging the Grid Control Management Framework](#)

## About Access and Identity Management

Access management is the means for controlling user access to enterprise resources. Access management products provide centralized, fine-grained access management for heterogeneous application environments, as well as out-of-the-box integration with Oracle products such as Oracle Portal, Oracle Collaboration Suite, and Oracle E-Business Suite.

Oracle Identity Management is a product set that allows enterprises to manage the end-to-end life cycle of user identities across all enterprise resources both within and beyond the firewall.

Automating user identity provisioning can reduce IT administration costs and improve security. Provisioning also plays an important role in regulatory compliance. Compliance initiatives focus on the enforcement of corporate policies as well as the demonstration of compliance with these standards. An enterprise identity management solution can provide a mechanism for implementing the user management aspects of a corporate policy, as well as a means to audit users and their access privileges.

Oracle Identity management products include the following:

- Access Manager
- Access Server
- Identity Server

### Access Manager

Oracle Access Manager, formerly known as Oracle COREid Access and Identity, provides Web-based identity administration and access control to Web applications and resources running in a heterogeneous environments. It provides the user and group management, delegated administration, password management and self-service functions necessary to manage large user populations in complex, directory-centric environments.

Access Manager supports all popular authentication methods including browser forms, digital certificates, and smart cards, and integrates seamlessly with most application servers and portals, including OracleAS 10g, Oracle WebLogic, IBM WebSphere, Vignette and others. User identities and credentials can be accessed from a number of repositories including Oracle Internet Directory, Microsoft Active Directory, and Sun Java System Directory. With Access Manager, user access policies can be defined and enforced with a high degree of granularity through centralized management.

### Access Server

The Access Server enables you to protect resources such as URLs and legacy, non-HTTP applications. It provides authentication and authorization services to enterprise applications. It uses the information stored by the Identity Server to control which users, groups, and organizations can access a resource. It stores information about configuration settings and security policies that control access to resources in a directory server that uses Oracle Access Manager-specific object classes. You can use the same directory to store the Access Server configuration settings, access policy data, and user data, or you can store this data on separate directory servers.

### Identity Server

The Identity Server is a set of applications that provide delegated administration, user self-service, and real-time change management. The Identity Server stores information

about users, groups, and organizations. For example, you can create, manage, and delete groups in the directory server. You can define a subscription policy for a group, including self-service with no approval needed, subscription with approvals, rule-based subscription, and no subscription allowed.

### **Identity Manager**

The Oracle Identity Manager platform automates user identity provisioning and deprovisioning and allows enterprises to manage the end-to-end lifecycle of user identities across all enterprise resources, both within and beyond the firewall. It provides an identity management platform that automates user provisioning, identity administration, and password management, wrapped in a comprehensive workflow engine.

Automating user identity provisioning can reduce IT administration costs and improve security. Provisioning also plays an important role in regulatory compliance. Key features of Oracle Identity Manager include password management, workflow and policy management, identity reconciliation, reporting and auditing, and extensibility through adapters.

Oracle Identity Manager also provides attestation support. Attestation is the process of having users or system managers confirm people's access rights on a periodic basis. Existing Sarbanes-Oxley requirements demand enterprises to perform attestation for all financially significant systems every three to six months. Identity Manager includes a highly flexible attestation solution to help enterprise customers meet these regulatory requirements in a cost-effective and timely manner. By setting up attestation processes in Identity Manager, enterprise customers can automate the process of generation, delivery, review, sign-off, delegation, tracking, and archiving of user access rights reports for reviewers on a scheduled or ad-hoc basis.

### **Identity Federation**

As more companies move their business processes to the Web, many organizations have a greater need to extend the boundaries of their enterprise to include partner applications. Federated identity management allows companies to operate independently and cooperate for business purposes by enabling cross-domain single sign-on and allowing companies to manage user identities and vouch for them as they access resources managed by another domain.

Oracle Identity Federation, formerly known as COREid Federation, provides a self-contained federation solution that combines the ease of use and portability of a standalone application with a scalable, standards-based proven interoperable architecture. It helps corporations securely link their business partners into a corporate portal or extranet while also increasing their compliance with privacy and security regulations. Identity Federation enables companies to manage multiple partners and choose from industry standard federated protocols. Identity Federation provides built-in integration with customer's identity management infrastructure (Oracle and non-Oracle) to deliver an end-to-end user experience, addressing scenarios like automatic registration, identity mapping, seamless access control navigation, and others.

## **Using Grid Control for Monitoring Identity Management Targets**

Enterprise Manager helps you monitor the availability and diagnose the health of Access, Identity, Identity Federation, and Identity Manager servers within your enterprise configuration. By deploying a Management Agent on each host, you can use Enterprise Manager to discover the Identity Management components on these

hosts, and automatically begin monitoring them using default monitoring levels, notification rules, and so on.

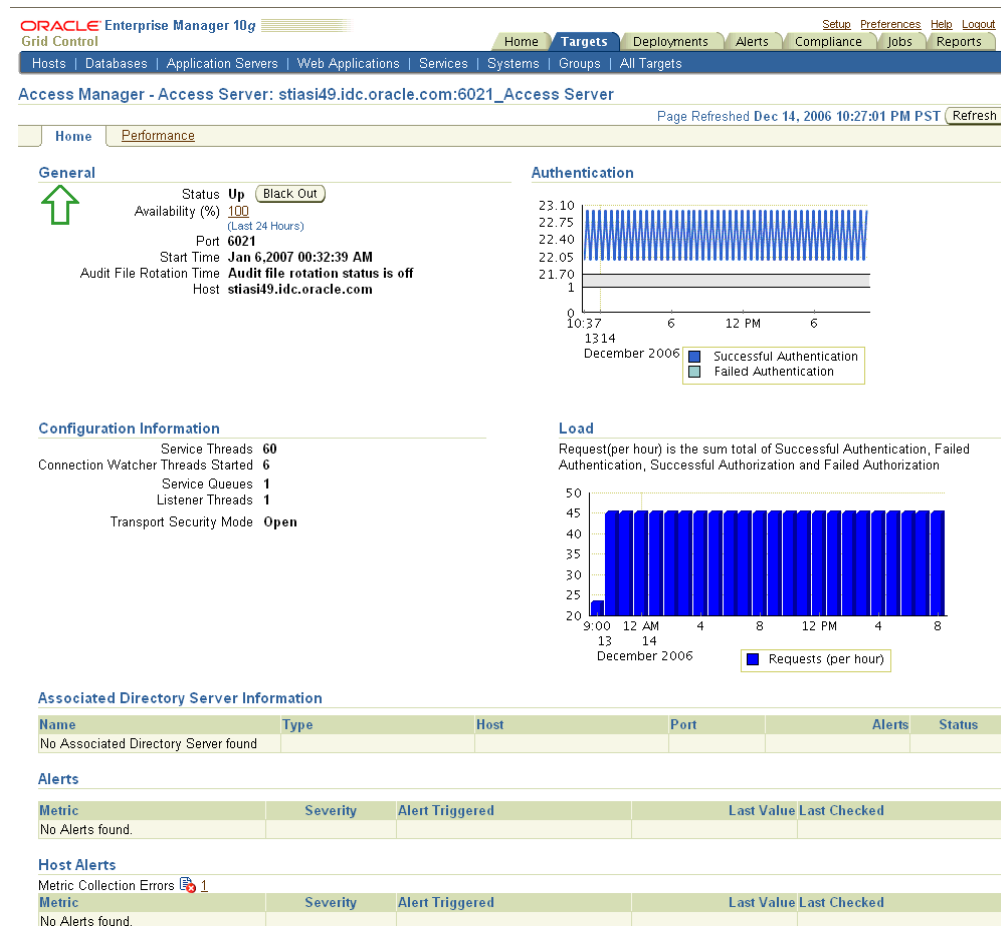
### Identity Component Server Home Page

All Identity Management targets, whether Access, Identity, Identity Federation, and Identity Manager have their own server home pages that provide easy access to key information required by the administrators. Each Identity Management Server home page provides the following information:

- Server status, responsiveness, and performance data
- Alerts and diagnostic drill-downs so you can identify and resolve problems quickly
- Resource usage for the server and its components
- Functionality to start, stop, and restart components in the case of locally-monitored Access and Identity servers
- Configuration parameters for Access and Identity servers

Figure 8–16 shows the Access Manager - Access Server home page.

**Figure 8–16 Access Manager - Access Server Home Page**



# Identity Management Systems

Identity Management services run on Identity Management systems defined in Grid Control. The system includes the software infrastructure components that the Identity services rely on. The system includes components such as databases, HTTP servers, OC4Js, and other servers.

The system is a collection of server targets that are grouped together in Grid Control to give you a view of the "data-center" components that comprise your Identity Management deployment. Identity Management Systems are created when Identity suite components are discovered using Grid Control. Grid Control also monitors the performance and availability of these components and provides a System Dashboard to view the health of the Identity Management system in a single window.

Figure 8-17 shows an Access Manager-Identity System home page:

Figure 8-17 Identity System Home Page

The screenshot displays the Oracle Enterprise Manager 10g Grid Control interface for the 'Access Manager - Identity System'. The top navigation bar includes 'Home', 'Targets', 'Deployments', 'Alerts', 'Compliance', 'Jobs', and 'Reports'. The main content area is divided into several sections:

- Status:** A green circle indicates 100% availability, with a 'Up(1)' indicator.
- ServiceTimePerRequest: Access Manager - Identity Server: Highest Average:** A line graph showing performance over time, with a legend for 'stias49.idc.oracle.com:6022\_Identity Server'.
- Services:** A table with columns for Name, Type, Status, Performance Alerts, Usage Alerts, and Policy Violations. It notes '(No data found.)'.
- Configuration Changes:** A section for configuration changes detected for the last 7 days.
- Load in the Last Collection Interval:** A table showing metrics: Request Processed (180), Successful Cache Flush Request (0), Cache Flush Request Failures (0), Email Notification Failures (0), and Total LDAP Sessions opened (9).
- Alerts:** A table with columns for Severity, Current, and Last 24 hours. It shows 0 alerts in all categories.
- Policy Violations:** A table with columns for Severity, Current, Last 24 Hours (Cleared, New), and Distinct Rules Violated. It shows '(No violations found)'. Below it is a section for Security Policy Violations, also showing '(No violations found)'. A link for 'Security At a Glance' is provided.
- Critical Patch Advisories for Oracle Homes:** A section for patch advisories.
- Related Links:** Links for Access Reports, Edit System, Target Properties, and Properties.

The footer contains copyright information for Oracle (1996, 2007) and a link to 'About Oracle Enterprise Manager'.

## Configuring Identity Management Systems

An Identity Management service is a logical target configured by Grid Control. You use Grid Control to step you through the process of configuring a web application

service for your Identity component instances. After you configure a service, that service is displayed on the Services page.

Critical application functions are defined and monitored as services in Grid Control. Each service is monitored by Grid Control beacons, which run service tests that simulate real user access to the service. Service availability and performance are monitored automatically, and problems are immediately reported to the administrator. By monitoring availability and performance of Identity Management services, you can identify and resolve user-visible problems more quickly and thus minimize the impact on users.

### **Creating an Access Service**

An access service allows you to perform service-level monitoring. If authentication or authorization services are not available, then the administrator should be notified about the service failures. Administrators will then be able to diagnose the cause of the problem using Root Cause Analysis.

1. Discover the Access Server target in Grid Control. For information about discovering this target, see the Grid Control online help. This will create the associated system.
2. In Grid Control, in the Services tab, select Add Generic Service and click Go. The Create Generic Service wizard is displayed.
3. In the Create Generic Service: General page, provide a name for the service and associate the Access system. Click Next.
4. In the Create Generic Service: Availability page, select Define availability based on Service Test. This allows you to create a web transaction. Click Next.
5. In the Create Generic Service: Service Test page, provide a name for the web transaction, select Record a Transaction and click Go.
6. In the Create Service Test page, in the Steps section, click Record.
7. In the Record Web Transaction page, click Start. A new browser window is displayed.
8. In the browser, enter the URL that is protected by the Access Server. In the Single Sign-On login page displayed, enter the username and password.
9. In the Record Web Transaction page, click Stop and then click Continue. The Create Web Applications: Service Test page is displayed again. Click Continue.
10. In the Create Generic Service: Beacons page, add a beacon to execute the web transaction you created. Click Next.
11. In the Create Generic Service: Performance Metrics page, click Next.
12. In the Create Generic Service: Usage Metrics page, click Next.
13. In the Create Generic Service: Review page, click Finish to create the Access service that will be monitored by the beacon test.

### **Creating an Identity Service**

Identity service allows the administrator to perform service-level monitoring of Access Manager Identity. If user management or group management services are not available, then the administrator should be notified about the service failures. Administrators will then be able to diagnose the cause of the problem using Root Cause Analysis.

1. Discover the Identity Server target in Grid Control. For information about discovering this target, see the Grid Control online help. This will create the associated system.
2. In Grid Control, in the Services tab, select Add Generic Service and click Go. The Create Generic Service wizard is displayed.
3. In the Create Generic Service: General page, provide a name for the service and associate the Identity system. Click Next.
4. In the Create Generic Service: Availability page, select Define availability based on Service Test. This allows you to create a web transaction. Click Next.
5. In the Create Generic Service: Service Test page, provide a name for the web transaction, select Record a Transaction and click Go.
6. In the Create Service Test page, in the Steps section, click Record.
7. In the Record Web Transaction page, click Start. A new browser window is displayed.
8. In the browser, enter the URL of the Identity system (<host:port/identity/oblix>) and then click the Identity System Console link. The login page is displayed.
9. In the Record Web Transaction page, click Stop and then click Continue. The Create Web Applications: Service Test page is displayed again. Click Continue.
10. In the Create Generic Service: Beacons page, add a beacon to execute the web transaction you created. Click Next.
11. In the Create Generic Service: Performance Metrics page, click Next.
12. In the Create Generic Service: Usage Metrics page, click Next.
13. In the Create Generic Service: Review page, click Finish to create the Identity service that will be monitored by the beacon test.

### **Creating an Identity Manager Service**

For information about creating a web application service for Identity Manager, see the "Creating Web Applications for Identity Manager" topic in online help.

### **Creating an Identity Federation Service**

You can create an Identity Federation service for service-level monitoring of Identity Federation. The procedure to create an Identity Federation service is similar to the procedure to create an Access or Identity service.

### **Monitoring Services**

Grid Control enables you to monitor all of your Identity Management services. Each service is monitored for performance, usage, and availability.

Each service has its own home page. The Service Home pages in Grid Control provide:

- Status, responsiveness, and performance data
- Resource usage data for the service
- Summary information such as status, performance alerts, usage alerts, and policy violations for the service's subcomponents, including other services and associated systems
- Links to home pages for the service's subcomponents

- Alerts and diagnostic drill-downs so that you can identify and resolve problems quickly
- Services Dashboard

The Services Dashboard provides a high-level view of the status, performance, and usage of each Identity Management target. Service-level compliance for various time periods are also included for each service on the dashboard. You can launch the dashboard directly from Identity system target home page. You can also publish the Services Dashboard so that it can be viewed by non-Enterprise Manager users. This allows you to provide a self-service status web page to your end users.
- Related Links to do the following:
  - View metrics for the service
  - View client configurations
  - Edit the service
  - View the service target's properties
  - Manage blackouts
  - View and manage metric thresholds and policies

**See Also:** [Chapter 6, "Service Management"](#)

## Identity Management Root Cause Analysis

Individual services in Identity Management are associated with critical system components. This allows Enterprise Manager to perform Root Cause Analysis down to the system level whenever a service outage is detected. When you are configuring an Identity Management service in Grid Control, you also mention the critical system components of this service. When an Identity Management service goes down, Enterprise Manager automatically performs a root cause analysis to determine which critical system component is responsible for this.

**See Also:** ["Diagnosing Service Problems"](#) in [Chapter 6, "Service Management"](#)

## Automated Identity Management and Alerts

Enterprise Manager automatically gathers and evaluates diagnostic information from Identity Management targets distributed across the enterprise. As with all targets managed by Enterprise Manager, an extensive number of Identity Management performance metrics are automatically monitored against predefined thresholds. Alerts are generated in Grid Control when metrics exceed these thresholds.

## Diagnosing Identity Management Performance and Availability Problems

You can use Grid Control to diagnose performance and availability problems with your Identity Management services. For example, if a service outage occurs, Root Cause Analysis will determine if the primary cause is an outage of a critical service or system component. If a service performance issue is found, an administrator can examine detailed metrics over time related to that service and any of the service or system components used by that service. When you suspect there is a problem with one or more server components in the Identity Management system, the system home pages provide metrics and charts for diagnosing the issue.



**See Also:** ["Monitoring Services"](#) and ["Diagnosing Service Problems"](#) in [Chapter 6, "Service Management"](#)

## Leveraging the Grid Control Management Framework

Grid Control includes many general features that are useful to an Identity Management administrator, including:

- **Job Automation:** You can use the Grid Control job system to schedule tasks you want to automate.
- **Policies:** You can utilize the policy framework to ensure your Identity Management infrastructure adheres to your site-specific standards.
- **Database and Application Server Management:** Using the single Grid Control console, you can also manage the specific databases and application servers in your Identity Management deployment if needed.
- **Extensions:** Grid Control also includes monitoring of key network components that may be part of your Identity Management deployment. You can also extend Grid Control to monitor other components that are not recognized out-of-box by Enterprise Manager.



---

---

## Host and Third-Party Target Management

Today's IT infrastructure has evolved to include a variety of server platforms from different vendors such as Red Hat, SUSE, Sun, IBM, HP, and Windows. A typical enterprise also contains hardware components and applications from a number of vendors to run services, provide storage, perform load balancing, and control IP traffic.

Out-of-box, Enterprise Manager allows you to centrally manage more than just your Oracle infrastructure:

- Monitor and maintain the operating system and hardware for the hosts running your Oracle software.
- Monitor IBM and Oracle application servers and clusters, F5 Server Load Balancers, Network Appliance Filers.

This chapter contains the following major sections:

- [Host Management](#)
- [Third-Party Target Management](#)

### Host Management

With the array of available server platforms, it becomes more and more difficult for system administrators to maintain the operating systems and hardware for the server, or host, on which Oracle software runs. With Grid Control, as soon as a Management Agent is deployed to the host, Grid Control automatically starts monitoring alert and configuration information for that machine. System administrators—and anyone who requires host information, such as DBAs managing databases on those hosts—benefit from Grid Control's host management feature.

---

---

**Note:** To access the host management pages in Grid Control, do the following:

1. Click the **Targets** tab.
2. Click the **Hosts** subtab.

This will display all the monitored hosts in your enterprise. Click the **Help** link on any page to access the Grid Control online help system.

---

---

Host management allows users to:

- Analyze performance trends for host hardware to predict future performance
- Measure service levels by monitoring host performance in real-time

- Validate software and hardware configuration across the enterprise

## Monitoring Hosts

Like for other managed targets, Grid Control's full suite of monitoring features, including alerts, custom notifications, blackouts, corrective actions, monitoring templates, and more, are available for hosts.

### Host Home Page

The Host Home page provides an overview of the status and vital statistics for each host that is part of the Grid Control environment.

Grid Control consolidates the relevant host information into a convenient single-screen Host Home page. You can see the availability, key configuration information, and outstanding alerts, as well as other pertinent information about the host. Convenient links allow you to view all the metrics collected for the host, change the thresholds as appropriate, or directly log in to the host to perform administrative actions.

Using the Host Home page, you can do the following:

- Drill down to view detailed metrics
- View operating system, hardware, and other configuration information for the host
- View policy violations and alerts for the host
- Analyze the job activity
- Determine whether there are outstanding patch advisories
- Determine the last security evaluation of the host
- Navigate to other pages to help you investigate the health of the host

Figure 9–1 Host Home Page

The screenshot displays the Oracle Enterprise Manager 10g interface for a specific host. The browser title is "Oracle Enterprise Manager (SYSMAN) - Host: dadvmm0630.us.oracle.com - Mozilla Firefox". The page has a navigation bar with tabs for Home, Targets, Deployments, Alerts, Compliance, Jobs, and Reports. The main content area is divided into several sections:

- General:** Shows the host status as "Up", last booted "Before 5 Days 11 Hours (Jan 10, 2009 2:34:33 AM (UTC-08:00))", 1 logon, and 100% availability. A "Black Out" button is visible.
- Configuration:** Lists system details such as Operating System (Enterprise Linux), Hardware Platform (i386), IP Address (10.229.149.34), CPU (1), Memory Size (4000 MB), and Local File Systems (96.46 GB).
- Alerts:** A table with columns for Metric Name, Severity, Message, Alert Triggered Value, and Last Checked. It currently shows "No alerts found".
- Policy Violations:** Indicates "Current 24 Distinct Policies Violated" and a "Compliance Score (%) 76".
- Security:** Shows the "Last Security Evaluation" on Jan 15, 2009, with a "Compliance Score (%) 68".
- Critical Patch Advisories:** Shows "Current 0" advisories and a note that patch information may be stale.
- Job Activity:** Displays counts for Problem Executions, Action Required Executions, Suspended Executions, Scheduled Executions, and Running Executions, all currently at 0.
- Related Links:** A collection of links for tasks like Access, Alert History, All Metrics, Blackouts, Deployments, Execute Host Command, Log File Alerts, Metric and Policy Settings, Monitoring Configuration, Net Services Administration, Open Telnet Session, Privilege Delegation Setting, Remote File Editor, Reports, Storage Details, Target Properties, and User-Defined Metrics.

The footer contains copyright information for Oracle and links to "About Oracle Enterprise Manager".

## Host Performance Page

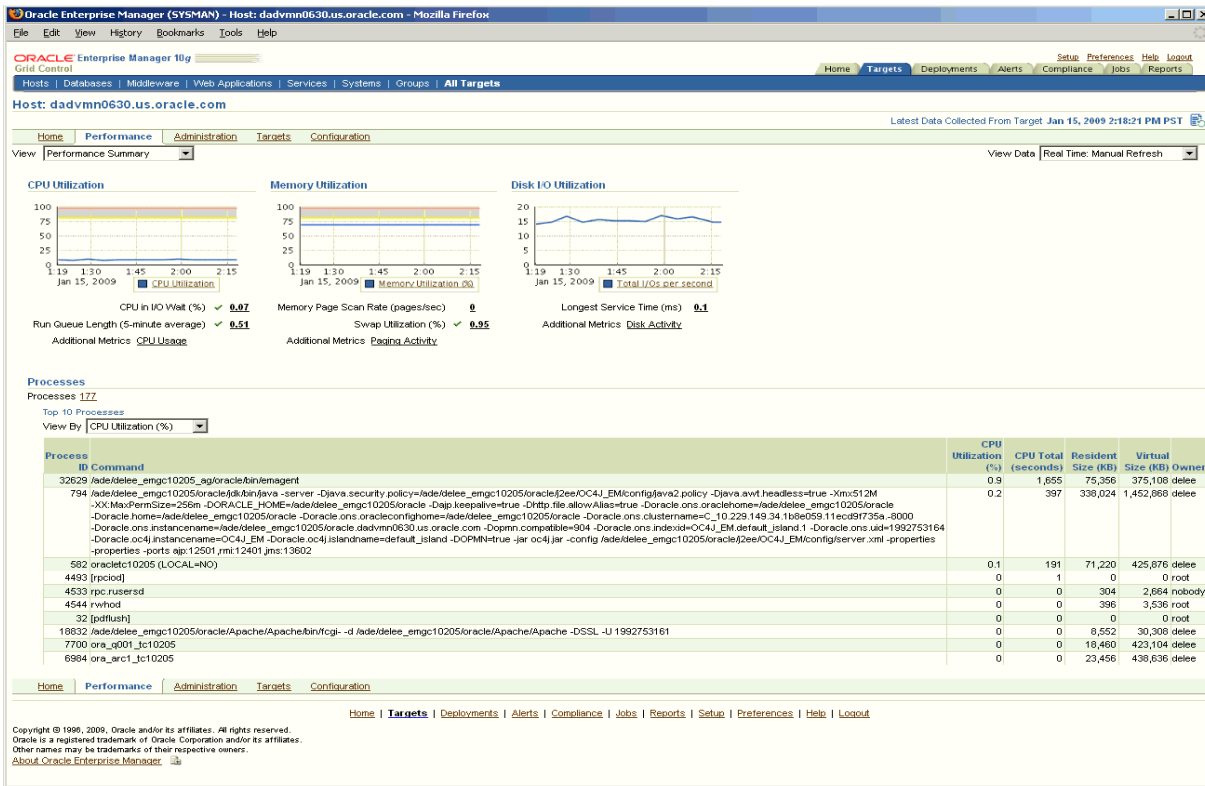
The Host Performance page provides an overview of the utilization statistics (CPU, Memory, Disk I/O, and Program Resource Utilization) for an individual host. Choose one of these options from the **View** menu. With this information, system administrators can determine whether resources need to be added or redistributed. You can also view the top processes consuming the most CPU or Memory and take appropriate action.

The performance metrics that are collected out-of-box for the host span several different categories: CPU Usage (including idle, wait, and user times), Disk Activity (including Average Disk I/O rate), and Network Interface Activity, among others. You can also view real-time metrics such as buffer activity.

Using the Performance page, you can:

- View various utilization charts by selecting an option in the **View** menu. Related metrics are also displayed under the charts. Click a metric link to view the metric in more detail.
- View the processes that are using the most CPU or memory resources.
- View performance data in real-time, or historical data over defined periods.

Figure 9–2 Host Performance Page



## Host Configuration Page

The Host Configuration page lists not only operating system and hardware configuration information, but all Oracle software and operating-system registered software on the host as well. You can also compare configurations for single or multiple hosts.

**See Also:** "Hardware and Software Configurations" in Chapter 4, "Enterprise Configuration Management"

## Operating System Details

Enterprise Manager surfaces key operating system-level information for hosts. You can obtain general information about the operating system, such as the distributor and the version. You can also drill-down for specific information such as system properties, file systems, and operating system-level packages. This information helps when troubleshooting performance problems that arise because of configuration issues.

**Note:** To access the Operating System Details page:

1. From the **Hosts** subtab, select a host target.
2. From the Host Home page, click the **Operating System Details** link under Configuration.

Figure 9–3 Host Operating System Details Page

ORACLE Enterprise Manager 10g  
Grid Control

Home | Targets | Deployments | Alerts | Policies | Jobs | Reports

Collaboration Suites | Hosts | Databases | Application Servers | Web Applications | Services | Systems | Groups | All Targets

Host [i1inrbn05.us.oracle.com](#) >

Operating System Details Data Collected Oct 13, 2005 5:37:46 AM PDT

Host **i1inrbn05.us.oracle.com** (History)

Operating System **Red Hat Enterprise Linux AS release 3 (Taroon Update 4) 2.4.21.27.ELsmp (32-bit)**

Vendor **Red Hat**

**General** | File Systems | Packages

Distributor Version **Red Hat Linux 3.2.3-47**  
Maximum Swap Space (MB) **8047.31**

**Operating System Properties**

Name	Source	Value
nameserver 1	/etc/resolv.conf	130.35.249.41
nameserver 2	/etc/resolv.conf	130.35.249.52
nameserver 3	/etc/resolv.conf	138.2.202.15
Maximum Swap Space (MB)	/sbin/sysctl	8047.31
abi.defhandler_coff	/sbin/sysctl	117440515
abi.defhandler_elf	/sbin/sysctl	0
abi.defhandler_lcall7	/sbin/sysctl	68157441
abi.defhandler_libcso	/sbin/sysctl	68157441
abi.fake_utsname	/sbin/sysctl	0
abi.trace	/sbin/sysctl	0
debug.kermltype	/sbin/sysctl	
debug.rpmarch	/sbin/sysctl	
dev.parpport.default.spintime	/sbin/sysctl	500
dev.parpport.default.timeslice	/sbin/sysctl	200
dev.raid.speed_limit_max	/sbin/sysctl	10000
dev.raid.speed_limit_min	/sbin/sysctl	100
dev.rtc.max-user-freq	/sbin/sysctl	64
fs.aio-max-pinned	/sbin/sysctl	393216
fs.aio-max-size	/sbin/sysctl	131072
fs.dir-notify-enable	/sbin/sysctl	1
fs.file-max	/sbin/sysctl	65536
fs.lease-break-time	/sbin/sysctl	45
fs.leases-enable	/sbin/sysctl	1
fs.nfs.nfs3_acl_max_entries	/sbin/sysctl	1024
fs.nfs.nlm_grace_period	/sbin/sysctl	0

Related Link [Search Operating System Property Settings on Hosts](#)

**General** | File Systems | Packages

Home | Targets | Deployments | Alerts | Policies | Jobs | Reports | Setup | Preferences | Help | Logout

The File Systems page includes information about the various mount points for the host, type of mount point, time of the mount, and the directory where the file system is mounted.

The Packages page lists all of the operating system packages installed on the host.

## Hardware Details

The Hardware Details page allows you to view the CPU, I/O, and Network Interfaces associated with the host. It also helps the user keep track of the hardware changes that occur over time. The type of operation (INSERT, UPDATE, or DELETE) and the category of the hardware that is updated are reflected on this page.

---

**Note:** To access the Hardware Details page:

1. From the **Hosts** subtab, select a host target.
  2. From the Host Home page, click the **Hardware Details** link under Configuration.
-

Figure 9–4 Host Hardware Details Page

Oracle Enterprise Manager (RPATTI) - Hardware Details - Mozilla Firefox

Oracle Enterprise Manager (SYSMAN) - Oracle Enterprise Manager (RPA...)

ORACLE Enterprise Manager 10g

Grid Control

Home Targets Deployments Alerts Compliance Jobs Reports

Hosts | Databases | Middleware | Web Applications | Services | Systems | Groups | All Targets

Host: stacv02.us.oracle.com >

Hardware Details Data Collected Jan 15, 2009 12:56:52 AM PST

System Configuration	Hostname	stacv02.us.oracle.com (stacv02)	Local Disk Capacity (GB)	71.39
Machine Architecture	System Configuration	i686	Clock Frequency (MHz)	1
Hardware Provider	Machine Architecture	GenuineIntel i686	Number of CPUs	1
Memory Size (MB)	Hardware Provider	Intel Based Hardware	Number of CPU boards	1
	Memory Size (MB)	2025	Number of IO devices	12

**CPUs**

CPU speed (MHZ)	Vendor	PROM Revision	ECACHE (MB)	CPU Implementation	Mask
2993	GenuineIntel	2	.5	Intel(R) Pentium(R) 4 CPU 3.00GHz	15

**IO Devices**

Name	Vendor	Bus Type	Frequency (MHZ)	PROM Revision
00:00.0 Host bridge	Intel Corporation 82875P/E7210 Memory Controller Hub	PCI	66	02
00:1d.0 USB Controller	Intel Corporation	PCI	66	02
00:1d.1 USB Controller	Intel Corporation	PCI	66	02
00:1d.2 USB Controller	Intel Corporation	PCI	66	02
00:1d.3 USB Controller	Intel Corporation	PCI	66	02
00:1d.7 USB Controller	Intel Corporation	PCI	66	02
00:1f.1 IDE interface	Intel Corporation	PCI	66	02
00:1f.2 IDE interface	Intel Corporation	PCI	66	02
00:1f.3 SMBus	Intel Corporation	PCI	66	02
02:01.0 Ethernet controller	Intel Corporation	PCI	66	
03:06.0 VGA compatible controller	Intel Corporation	PCI	66	27
03:08.0 Ethernet controller	Intel Corporation	PCI	66	01

**Network Interfaces**

Name	INET Address	Maximum Transfer Unit	Broadcast Address	Mask	Flags	MAC Address	Hostname Aliases
eth0		1500			BROADCAST,MULTICAST	00:0C:F1:AD:12:E6	
eth1	140.87.3.110	1500	140.87.3.255	255.255.252.0	BROADCAST,MULTICAST,RUNNING,UP	00:0C:F1:AD:12:E5	stacv02,stacv02.us.oracle.com
lo	127.0.0.1	16436		255.0.0.0	LOOPBACK,RUNNING,UP		localhost.localdomain,localhost
si0		1480			NOARP		

TIP Some Information may not be available depending upon the Hardware platform.

Home | Targets | Deployments | Alerts | Compliance | Jobs | Reports | Setup | Preferences | Help | Logout

Copyright © 1996, 2009, Oracle and/or its affiliates. All rights reserved.  
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.  
Other names may be trademarks of their respective owners.  
[About Oracle Enterprise Manager](#)

## Log File Monitoring

Grid Control monitors the log files for the occurrence of operator-specified patterns. Use this facility to monitor abnormal conditions recorded in the log files present on the host.

Log files are periodically scanned for the occurrence of desired patterns and an alert is raised when the pattern occurs during a given scan. During a scan, new content created since the last scan is searched for the occurrence of the desired patterns. Use this page to view, clear, and purge open alerts generated during log file monitoring.

---

**Note:** To access the Log File Alerts page:

1. From the **Hosts** subtab, select a host target.
  2. From the Host Home page, click the **Log File Alerts** link under Related Links.
- 

**See Also:** "Configuring Generic Log File Monitoring Criteria" in Grid Control online help

## Program Resource Utilization Monitoring

Enterprise Manager allows you to track resource use for programs and users. For example, you can track CPU usage by user, by program, or a combination of the two.

The Program Resource Utilization page provides a quick glimpse of the programs being monitored on this host. With this information, you can see trends in resource usage for:

- A specific program or set of programs



- A specific user or set of users
- A combination of programs and users

---

---

**Note:** To access the Program Resource Utilization page:

1. From the **Hosts** subtab, select a host target.
2. From the Host Home page, go to the **Performance** page.
3. From the Host Performance page, select **Program Resource Utilization** from the View menu.

You can also access the page from the All Metrics page for that host.

---

---

**See Also:** "Configuring Program Resource Utilization Monitoring Criteria" in Grid Control online help

### File and Directory Monitoring

Enterprise Manager monitors the files and directories for the operator-specified criteria on hosts running various versions of the UNIX operating system. The operator should configure the criteria for monitoring the desired files and directories.

---

---

**Note:** To access the File Directory Monitoring page:

1. From the **Hosts** subtab, select a host target.
  2. From the Host Home page, click the **All Metrics** link under Related Links.
  3. From the All Metrics page, select **File and Directory Monitoring** from the list of metrics.
- 
- 

**See Also:** "Configuring File and Directory Monitoring Criteria" in Grid Control online help

### Hardware Monitoring for Dell PowerEdge Systems

Hardware-specific monitoring is available out-of-box for Dell PowerEdge Linux hosts with Enterprise Manager. The following hardware health statistics can be monitored as part of the Dell PowerEdge Linux host target:

- Processor Status
- Memory Status
- PCI Device Status
- Power Supply Status
- System BIOS Status
- Fan Status
- Remote Access Card Status
- Temperature Probe Status

---

---

**Note:** To access the Hardware Monitoring page for Dell PowerEdge Systems:

1. From the **Hosts** subtab, select a host target.
  2. From the Host Home page, click the **All Metrics** link under Related Links.
  3. From the All Metrics page, select the metric you want to monitor from the list of metrics.
- 
- 

**See Also:** "Enabling Hardware Monitoring for Dell PowerEdge Linux Hosts" in Grid Control online help

## Storage Resource Tracking

Tracking the storage resource allocation and usage is essential to large IT departments. Unallocated and underutilized storage can be put to better use. Historical trends at a business entity-level enable you to plan for future growth.

Storage Details are relevant to Enterprise Manager targets that are associated with one or more hosts. In particular:

- Summary attributes presented are rolled up for one or multiple associated hosts.
- Storage Details helps in tracking storage resource allocation and usage for one host or a group of hosts, and tracks historical usage trends.
- A host is associated with a group either through explicit membership, or implicit membership inherited through a group member target.

---

---

**Note:** To access the Storage Details page:

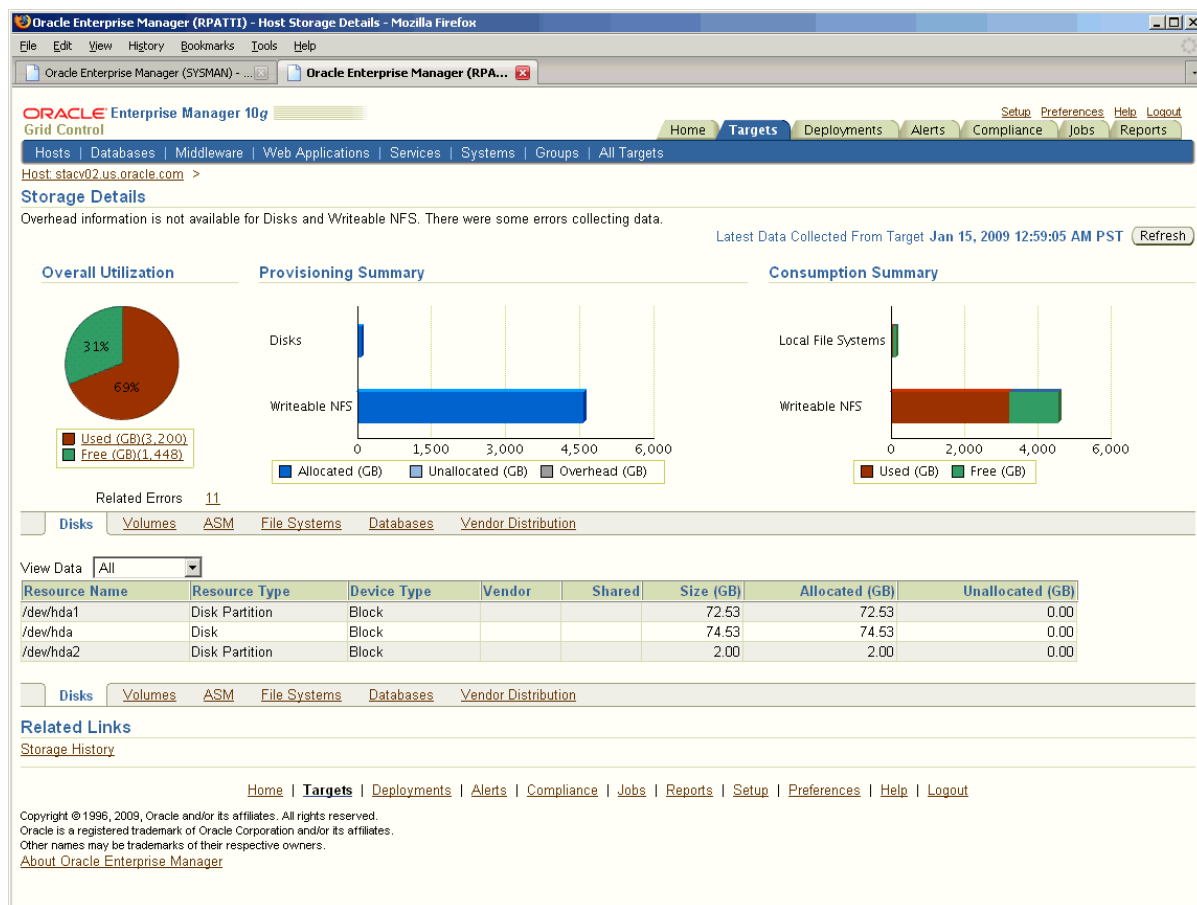
1. From the **Hosts** subtab, select a host target.
  2. From the Host Home page, click the **Storage Details** link under Related Links.
- 
- 

The Storage Details page displays the following charts:

- **Overall Utilization:** Shows summary attributes that provide a system-level view of storage resource utilization.
- **Provisioning Summary:** Shows allocation-related summary attributes for File Systems, ASM, Volumes, and Disks for the associated hosts.
- **Consumption Summary:** Shows usage-related summary attributes for Databases and File Systems.

The Disks option on the Storage Details page shows the allocated and unallocated storage for all the disks and disk partitions on a host. All disks are listed, including virtual disks from external storage systems such as EMC Storage Array.

Figure 9–5 Host Storage Details Page



## Host Administration

In addition to the general Grid Control features and tools described in the "Overview of Enterprise Manager" chapter, you can administer your hosts using the Host Administration tab and Grid Control's tools designed for host management.

### Linux Host Administration

Linux host administration features for RedHat (RHEL4) and SuSE Linux (sles9) provide you the ability to administer the most important host functions easily and quickly. Operators with the appropriate privileges can set up and administer the following OS features:

- System Services
  - Start and stop individual services — Facilitates upgrading and troubleshooting applications and databases
  - Configure services to run at boot time — Ensures that all your applications and dependent services are available after every reboot
  - Assign system services to run levels — Assures appropriate behaviors to assist in maintenance and trouble-shooting
- Network Setup

- NFS Client Configuration — Allows you to configure access to storage resources
- Network Card Configuration — Enables you to rapidly establish connectivity to the network to get your application service running quickly
- View and edit host name lookup tables (/etc/hosts) — Allows you to set-up and edit host names to identify new services to the network

---

---

**Note:** To access the Host Administration page:

1. From the **Hosts** subtab, select a Linux host target.
  2. From the Host Home page, click **Administration**.
- 
- 

### Execute Host Command

With the appropriate privileges, you can execute non-interactive host commands remotely over the network, as well as view and edit text files. The Execute Host Command page enables you to type operating system commands against one or more hosts, or all the hosts in a group, enabling you to perform multiple administrative operations at the same time.

---

---

**Note:** To access the Execute Host Command page:

1. From the **Hosts** subtab, select a host target.
  2. From the Host Home page, click the **Execute Host Command** link under Related Links.
- 
- 

### Job System

You can leverage the Job System to run operating system commands for your host. From the Execute Host Command page, use the "Load from Job Library" option to search the Job Library for existing jobs that you can reuse.

From the Load from Job Library page, click the icon in the **Load** column of any row to return to the Execute Host Command page, loading the complete context of the library job in that row. The complete context includes the host command, OS script, targets, and credentials.

---

---

**Note:** To access the Load from Job Library page:

1. From the Execute Host Command page, click the **Switch to Multiple Target Mode** link.
  2. Click the **Load from Job Library** button.
- 
- 

### Remote File Editor

The Remote File Editor page enables you to view, edit, copy, and save text files present on the remote host target. For example, using this utility, you can update the contents of configuration files on the remote host.

---

---

**Note:** To access the Remote File Editor page:

1. From the **Hosts** subtab, select a host target.
  2. From the Host Home page, click the **Remote File Editor** link under Related Links.
- 
-

## Managing Host Deployments

Use Grid Control's deployment management tools to:

- View and compare host configurations
- Perform automated patching for your hosts
- Deploy new hosts using the provisioning facility

**See Also:** ["Linux Host Patching" in Chapter 13, "Lifecycle Management"](#)

## Third-Party Target Management

Grid Control enables you to manage hardware and software from certain third-party vendors out-of-box. Most tools, functions, and features available for monitoring and administering Oracle targets can be used to manage other targets in your data center.

---

---

**Note:** Use Grid Control to manage your complete data center by developing Management Plug-ins for your managed targets. Refer to *Enterprise Manager Extensibility* for more information.

---

---

## Support for Third-Party Application Servers

With tools and features similar to managing Oracle Application Server targets, Grid Control allows you to monitor the following third-party application server targets, as well as the applications deployed on the servers:

- Oracle WebLogic Domains, Clusters, and Managed Servers
- IBM WebSphere Cells, Clusters, and Application Servers
- JBoss Partitions and JBoss Application Servers
- Microsoft Exchange Organizations, Routing Groups, and Servers

Oracle WebLogic Domains, IBM WebSphere Cells, and Microsoft Exchange Organizations represent logically-related groups of application server resources. Grid Control monitors them as composite targets, containing rolled-up information on the member targets that make up the group.

All the resources in a Oracle WebLogic Domain are centrally managed by the WebLogic Administration Server, and the resources in an IBM WebSphere Cell are managed by the Deployment Manager. The JMX fetchlets specific to each product environment communicate with the Oracle WebLogic Administration Server or IBM WebSphere Deployment Manager to collect the metric information for the individual application servers that make up the Domain or Cell, respectively. There is no need to deploy Oracle Management Agents to each application server node; as long as a Management Agent is locally or remotely deployed to the host on which the Administration Server or Deployment Manager resides, Grid Control can manage those targets.

The member targets that make up a Oracle WebLogic Domain are clusters and managed servers, and the member targets that make up an IBM WebSphere Cell are clusters and application servers. As for Microsoft Exchange Organization, the member targets are individual Microsoft Exchange Servers.

Oracle WebLogic Cluster, IBM WebSphere Cluster, JBoss Partition, and Microsoft Exchange Routing Group provide a high-availability application server environment. Grid Control monitors them as composite targets, containing rolled-up information on

the member targets that make up the group. The member targets for each of these targets are the application servers grouped under them.

Oracle WebLogic Managed Servers, IBM WebSphere Application Servers, JBoss Application Servers, and Microsoft Exchange Servers are used to deploy applications and other resources. They are all monitored as individual targets, much like Oracle Application Servers.

Grid Control also provides configuration management features for these third-party application servers. Using these features, you can view the configurations that were last collected, save configurations so that they can be referenced in the future for possible analysis or comparison, and compare configurations to examine specific configuration variations between two instances. You can also perform some out-of-box searches, for example, you can search one or more J2EE applications deployed in your third-party application server instances.

However, note that all these third-party application servers have to be manually added to Grid Control for monitoring purpose. Unlike Oracle Application Servers, these third-party application servers are not automatically discovered by Grid Control. You can use the Application Servers tab of the Grid Control console to add these targets.

Some of the features available for these third-party application server targets are:

- **Robust monitoring:** Automatic performance and availability monitoring with Oracle recommended settings, out-of-box notifications for critical alerts, J2EE application monitoring for clusters, J2EE services monitoring for JBoss application servers, historical collections, blackouts, monitoring of application-defined MBeans, and more.
- **Service management:** Transaction performance monitoring, end-user performance monitoring, system monitoring, dashboards, Root Cause Analysis, and more.
- **Group management:** Group-level roll-ups for monitoring and task automation, monitoring templates.
- **Reporting:** Out-of-box and user-defined reports for JBoss Application Server.
- **Policy violations:** Out-of-box and user-defined policies for Oracle WebLogic Managed Server.
- **Administration:** Process control for starting, stopping, and restarting Oracle WebLogic Managed Server and IBM WebSphere Application Server. Also, control for starting and stopping Channels of IBM WebSphere MQ targets.
- **Configuration Management:** Viewing, comparing, and searching configuration information related to these application servers.

**See Also:** "Managing Oracle WebLogic Servers", "Managing IBM WebSphere Application Server", "Managing JBoss Application Server", and "Managing Microsoft Exchange Server" in Grid Control online help.

## Support for Third-Party Plug-Ins

Much like the support offered for third-party application servers, Grid Control also supports third-party plug-ins like IBM WebSphere MQ Queue Manager and IBM WebSphere MQ Cluster.

IBM WebSphere MQ Cluster is a composite target that groups the IBM WebSphere MQ Queue Managers, and provides a high-availability environment. Grid Control

monitors them as logical targets and provides rolled-up information about the queue managers that make up the cluster.

IBM WebSphere MQ Queue Manager is monitored like any other application server target. You can monitor their status, be notified for critical alerts, access out-of-box reports and policies, and view a list of metrics and gauge their performance.

The other logical objects of IBM WebSphere MQ Queue Manager for which Grid Control collects metrics are Queues and Channels. For each of these logical targets, Grid Control provides a home page to help you monitor their status and performance.

**See Also:** "Managing IBM WebSphere MQ Queue Manager" and "Managing IBM WebSphere MQ Cluster" helpsets in Grid Control online help

## Network Appliance Filer

When you use Oracle Enterprise Manager to manage your Oracle environment, you can monitor not only your databases and application servers, but also supporting hardware, such as IP traffic controllers and storage devices.

Network Appliance Filers configured to use the ONTAP operating software (release 6.0 or later) are one type of storage device that can be managed out of the box. Network storage devices can be an important part of your Oracle environment. Use them to store and access large amounts of database and application data.

Similar to other managed targets, Grid Control's full suite of monitoring features are available for Network Appliance Filers. Key aspects of Network Appliance Filer monitoring are displayed as part of the Network Appliance-specific target pages, and include the following:

- Availability
  - Filer Target
  - Cluster Filer's Partner Status
- Storage Space Allocation and Usage Monitoring
  - Traditional and Flexible Volumes
  - Aggregates, Qtrees, and Snapshots
  - System Level Storage Capacity Summaries
- Performance Monitoring
  - CPU Utilization and NFS Statistics
  - NFS/CIFS Calls
  - Disk, Network, and Network Interface I/O
- Health Monitoring
  - Power Supplies, Fans, and NVRAM Batteries
  - Temperature and NFS/CIFS Bad Calls
  - Cluster Partner/Interconnect Status
  - Network Interface Status/Errors/Discards
  - Disk Failures and Spare Disks

**See Also:** "Overview of Adding a Network Appliance Filer" in Grid Control online help

## F5 Load Balancer Switches

In addition to monitoring the performance of your database, the effectiveness of your mid-tier applications, and the response time of your Web pages, Grid Control helps you monitor load balancer switches. Load balancer switches are an important part of any complex Web application environment, because they distribute user requests to the Web servers.

F5 BIG-IP Load Balancers from F5 Networks running OS version 4.2 PTF-06, or later, can be managed out of the box from Grid Control. The user interface allows you to navigate through the relationships that exist among the entities present in the switch.

As with other managed targets, Grid Control's full suite of monitoring features are available for F5 Load Balancer Switches. Key aspects of F5 Load Balancer Switch monitoring are displayed on the Load Balancer Switch-specific target pages, and include the following:

- Availability
  - Load Balancer Switch Target
  - Virtual and Real Server Status
- Performance Monitoring
  - Switch Level — CPU and Memory Utilization
  - Switch, Virtual Server, Server Group and Real Server Level — Bits In and Bits Out Rates, Connections Per Second
  - Physical Interface Level — Bits In and Bits Out Rates

**See Also:** "Overview of Adding a Load Balancer Switch" in Grid Control online help



---

---

## Oracle Collaboration Suite Management

Oracle Collaboration Suite provides capabilities—mail, calendar, web conferencing, and content services—that are critical to the productivity of an organization. Users access Collaboration Suite services every day to communicate with colleagues, host remote meetings, and share information; therefore, it is important that Collaboration Suite administrators have reliable, comprehensive tools to effectively manage the performance and availability of their Collaboration Suite deployments.

Oracle Enterprise Manager provides the monitoring and administration capabilities administrators need to ensure that the performance and availability of Collaboration Suite services meet customer expectations and service level goals.

This chapter describes how you can use Grid Control to manage your Oracle Collaboration Suite instances and components.

This chapter contains the following sections:

- [Visualizing Your Collaboration Suite Deployment](#)
- [Automated Collaboration Suite Monitoring and Alerts](#)
- [Oracle Collaboration Suite Root Cause Analysis](#)
- [Diagnosing Oracle Collaboration Suite Performance and Availability Problems](#)
- [Collecting and Using Collaboration Suite Client Configuration Data](#)
- [Oracle Collaboration Suite Reports](#)
- [Administration for Oracle Collaboration Suite Instances](#)

### Visualizing Your Collaboration Suite Deployment

Oracle Enterprise Manager utilizes the following Enterprise Manager concepts to visualize your Collaboration Suite deployment:

- [Collaboration Suite Services](#)
- [Collaboration Suite Systems](#)

### Collaboration Suite Services

Critical application functions are defined and monitored as services in Grid Control. Each Collaboration Suite component (such as Mail, Calendar, Content Services, and so on) is a service. In addition, each component is made up of a set of services.

For example, the Mail service includes services that represent IMAP, SMTP, and the Web Mail application. By breaking down the Mail service into individual sub-services, you can monitor specific Mail functionality separately. You can then, for example,

distinguish between all Mail services being unavailable versus just the Web Mail application being inaccessible to users.

Because Collaboration Suite provides critical end-user functionality, it is useful to treat these functions as services. Each service is monitored by Grid Control beacons, which run service tests that simulate real user access to the service. Service availability and performance are monitored automatically, and problems are immediately reported to the administrator. By monitoring availability and performance of Collaboration Suite services, you can identify and resolve user-visible problems more quickly and thus minimize the impact on users.

The hierarchy of services that represent Collaboration Suite is used to visualize a Collaboration Suite deployment in Grid Control. The Collaboration Suite target subtab shows Collaboration Suite services in a hierarchical listing along with critical performance, usage, and availability data for each. You can also see an overview of Collaboration Suite performance, usage, and availability in the Services Dashboard. The dashboard also shows service-level compliance data. Finally, the service topology is a graphical tree view of your Collaboration Suite deployment.

Oracle Enterprise Manager also provides the following features to help you visualize, monitor, and manage your Collaboration Suite instances and components:

- **Collaboration Suites Tab**

After you add the Collaboration Suites tab in Grid Control, you can configure a Collaboration Suite service for each of your Collaboration Suite instances, and manage all the instances from the Collaboration Suites tab.

**See Also:** "Adding the Collaboration Suites Tab" in the Grid Control online help

- **Configuring Collaboration Suite Services**

A Collaboration Suite service is a logical target configured by Grid Control that models the hierarchy of a Collaboration Suite instance and its components. You use Grid Control to step you through the process of configuring a Collaboration Suite service for each of your Collaboration Suite instances. After you configure a Collaboration Suite service, that service is displayed on the Collaboration Suite Services page.

When you configure a Collaboration Suite service for a Collaboration Suite instance, Grid Control also configures services for the Collaboration Suite components.

- **Collaboration Suite Services Page**

The Collaboration Suite Services page provides you with a graphical, tree view of Collaboration Suite services. On this page, you can see the relationship between the various services that make up your Collaboration Suite deployment. This page also provides summary information about the availability, performance, and usage of each of the Collaboration Suite instances for which you have configured a Collaboration Suite service. From the Collaboration Suite Services page, you can drill down to more detailed information about your Collaboration Suite instances and components.

- **Collaboration Suite Services Dashboard**

The Services Dashboard provides a high-level view of the status, performance, and usage of each Collaboration Suite component. Service-level compliance for various time periods are also included for each service on the dashboard. You can launch the dashboard directly from the Collaboration Suite home page. You can

also publish the Services Dashboard so that it can be viewed by non-Enterprise Manager users. This allows you to provide a self-service status web page to your end users.

**Figure 10–1 Collaboration Suite Service Home Page**

The screenshot shows the Oracle Enterprise Manager Grid Control interface for the 'Voicemail and Fax Service'. The page is titled 'Voicemail and Fax Service : Voicemail and Fax Service 3' and includes a navigation menu with options like Home, Charts, and Topology. The main content area is divided into several sections:

- General:** Shows the service status as 'Up' (with a 'Black Out' button), 'Up Since Oct 21, 2005 3:15:43 PM', 'Availability (%) 100 (Last 24 Hours)', 'Performance' (green checkmark), 'Usage' (green checkmark), 'Actual Service Level (%) 100.0000 (Last 24 Hours)', and 'Expected Service Level (%) 85.0000'.
- Subservices:** A table listing various subservices with columns for Name, Type, Status, Performance Alerts, Usage Alerts, Policy Violations, Name, and Key Components (Status, Alerts).
- Alerts:** A section with a 'View' dropdown set to 'All Alerts' and a table with columns for Target Name, Target Type, Severity, Alert Triggered, and Message. It currently shows 'No alerts found'.
- Related Links:** A section with links for Access, Edit Service, Reports, All Metrics, Edit Service Level Rule, Root Cause Analysis Configuration, Blackouts, Metric and Policy Settings, and Target Properties.

### Service Monitoring

Grid Control enables you to monitor all of your Collaboration Suite services and their component services. Each service is monitored for performance, usage, and availability.

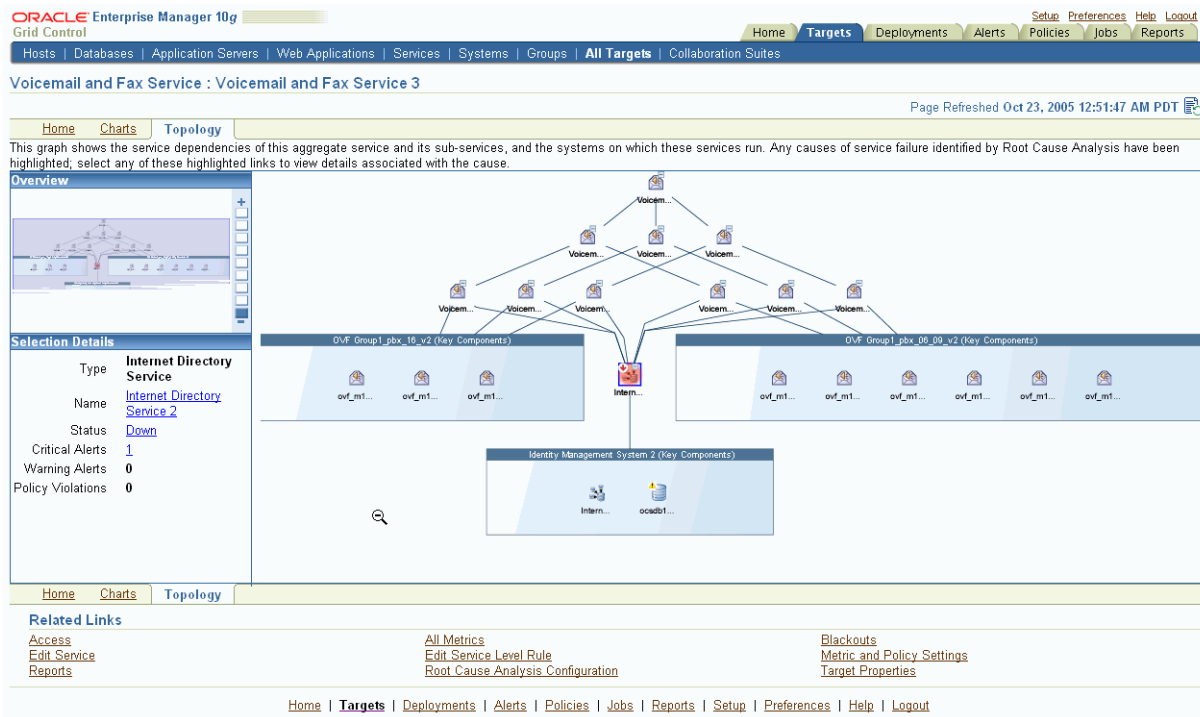
Each service has its own home page. The Service Home pages in Grid Control provide:

- Status, responsiveness, and performance data
- Resource usage data for the service
- Summary information such as status, performance alerts, usage alerts, and policy violations for the service's subcomponents, including other services and associated systems
- Links to home pages for the service's subcomponents
- Alerts and diagnostic drill-downs so that you can identify and resolve problems quickly
- The service's Topology Viewer
- Access to the Services Dashboard
- Related Links to do the following:

- View metrics for the service
- View client configurations
- Manage reports
- Edit the service
- View the service target's properties
- Manage blackouts
- View and manage metric thresholds and policies

**See Also:** [Chapter 6, "Service Management"](#)

**Figure 10–2 Collaboration Suite Service Topology Page**



## Collaboration Suite Systems

Collaboration Suite services run on a Collaboration Suite system defined in Grid Control. The system includes the software infrastructure components that the Collaboration Suite services rely on. The system includes components such as databases, HTTP servers, OC4Js, and other servers (such as IMAP or Calendar servers) used by Collaboration Suite.

The system is a collection of server targets that are grouped together in Grid Control to give you a view of the "data-center" components that comprise your Collaboration Suite. Grid Control also monitors the performance and availability of these components and provides a System Dashboard to view the health of the Collaboration Suite system in a single window.

## Automated Collaboration Suite Monitoring and Alerts

Enterprise Manager automatically gathers and evaluates diagnostic information from Collaboration Suite targets distributed across the enterprise. As with all targets managed by Enterprise Manager, an extensive number of Collaboration Suite performance metrics are automatically monitored against predefined thresholds. Alerts are generated in Grid Control when metrics exceed these thresholds.

## Oracle Collaboration Suite Root Cause Analysis

Individual services in Collaboration Suite are associated with the critical service and system components they rely on. This allows Enterprise Manager to perform Root Cause Analysis down to the system level whenever a service outage is detected. When you are configuring a Collaboration Suite service or component service in Grid Control, some of the UI pages give you the option of providing information about beacons, databases, and application servers that are critical for the component. If you do specify this information, Grid Control will be able to take these beacons, databases, and application servers into account when performing a Root Cause Analysis for the service.

**See Also:** ["Diagnosing Service Problems" in Chapter 6, "Service Management"](#)

## Diagnosing Oracle Collaboration Suite Performance and Availability Problems

You can use Grid Control to diagnose performance and availability problems with your Oracle Collaboration Suite services. For example, if a service outage occurs, Root Cause Analysis will determine if the primary cause is an outage of a critical service or system component. If a service performance issue is found, an administrator can examine detailed metrics over time related to that service and any of the service or system components used by that service. When you suspect there is a problem with one or more server components in the Collaboration Suite system, the system home pages provide metrics and charts for diagnosing the issue.

**See Also:** ["Monitoring Services"](#) and ["Diagnosing Service Problems" in Chapter 6, "Service Management"](#)

## Collecting and Using Collaboration Suite Client Configuration Data

The Oracle Collaboration Suite extensions to the basic Client System Analyzer (CSA) application enable Collaboration Suite-specific information to be collected from Windows client machines. After CSA for Collaboration Suite collects the information, it uploads it to the Grid Control Management Repository.

**See Also:** ["Client Configurations" in Chapter 4, "Enterprise Configuration Management"](#)

Collaboration Suite administrators and help desk personnel assisting Collaboration Suite users can use the collected information to identify clients that are not properly configured to use Oracle Collaboration Suite.

After installing and configuring the Collaboration Suite enhancements to CSA, the Collaboration Suite administrator can provide users with the URL for the CSA for the Collaboration Suite application. When users access the URL, they can give permission

for the CSA applet to be downloaded to their Windows client machine. After it is downloaded, the applet collects Collaboration Suite-specific information on the Windows client machine, including:

- Operating system user privileges. Only a user with administrator privilege can install Collaboration Suite Connector for Microsoft Office. CSA will check whether a user is an administrator.
- Type, version, and installation location of the mail client. CSA can identify Microsoft Office, Netscape, Mozilla, and Outlook Express mail clients on the client machine. If there are multiple mail clients installed, information on all of them will be collected.
- IMAP server and SMTP server settings.
- Execution mode for Outlook Express 98 or 2000 mail clients. Oracle Connector for Microsoft Outlook Express only works if Outlook Express is configured in Workgroup mode. In Outlook Express 2002 and later, the execution mode is always set to Workgroup.
- Type of Oracle Calendar client installed. If the Oracle Calendar thick client is installed, information about the version and install location of Oracle Calendar is collected.
- Browser versions and minimum recommended versions for optimally using various Collaboration Suite components, such as Web Access Client and Web Conferencing.
- Latency between the client machine and mail servers (IMAP, NNTP, POP, and SMTP).

Users who use CSA will see the summary report the applet displays after it collects the Collaboration Suite-specific data. The report identifies the areas in which the client does not meet Oracle Collaboration Suite requirements.

Administrators can use Grid Control to view the Oracle Collaboration Suite-specific client configuration information collected by CSA by navigating to the Collaboration Suite Service page, then clicking the **Client Configurations** link in the Related Topics section.

**See Also:** "Viewing Collaboration Suite Client Configuration Information" in the Grid Control online help

## Oracle Collaboration Suite Reports

You can define reports for Oracle Collaboration Suite services and systems, and also for Collaboration Suite component services, systems, and redundancy groups. For this release, you can view the following new reports for mail systems:

- Mail Queue Information Report
- Mail User Information Report
- Mail Message Information Report
- Mail Application-Tier Host Information Report
- Mail Detail Process Information Report
- Search Mail Message Report
- Unprocessed Mail Messages Report
- Mail Quota Usage Report

## Leveraging the Grid Control Management Framework

Grid Control includes many general features that are useful to a Collaboration Suite administrator, including:

- **Job Automation:** You can use the Grid Control job system to schedule tasks you want to automate.
- **Patching:** The configuration management capabilities in Grid Control allow you to manage and apply patches to your Collaboration Suite installations. You are also alerted automatically when critical patch updates are released.
- **Policies:** You can utilize the policy framework to ensure your Collaboration Suite infrastructure adheres to your site-specific standards.
- **Database and Application Server Management:** Using the single Grid Control console, you can also manage the specific databases and application servers in your Collaboration Suite deployment if needed.
- **Extensions:** Grid Control also includes monitoring of key network components that may be part of your Collaboration Suite deployment. You can also extend Grid Control to monitor other components that are not recognized out-of-box by Enterprise Manager.

## Administration for Oracle Collaboration Suite Instances

When you install an Oracle Collaboration Suite instance, you automatically get Oracle Application Server Control for Collaboration Suite to manage that instance. Each installed Collaboration Suite instance has its own Application Server Control for Collaboration Suite.

Application Server Control Console for Collaboration Suite is the Enterprise Manager Web-based application for managing Oracle Collaboration Suite releases 9.0.4 and 10g Release 1 (10.1). Application Server Control for Collaboration Suite provides Web-based management tools designed specifically for Collaboration Suite.

You can use Application Server Control for Collaboration Suite to monitor and administer a single Collaboration Suite instance, a group of Collaboration Suite instances, or individual components for a Collaboration Suite instance. You can also deploy applications, monitor real-time performance, manage security, and configure the components of your Collaboration Suite instance.

Application Server Control for Collaboration Suite relies on various underlying technologies to discover, monitor, and administer the Collaboration Suite environment.

Application Server Control for Collaboration Suite consists of Application Server Control Console for Collaboration Suite and its underlying technologies:

- Oracle Dynamic Monitoring Service (DMS)
- Oracle Process Management Notification (OPMN)
- Distributed Configuration Management (DCM)
- A local version of the Oracle Management Agent specifically designed to gather monitoring data for the Application Server for Collaboration Suite Control.

For additional management functionality, for example, Service Level Management, deployments, historical data collections for performance trending alerts, and so on, you can use Enterprise Manager.

**See Also:** *Oracle Collaboration Suite 10gR1 Administrator's Guide*

While Application Server Control for Collaboration Suite provides standalone management for a Collaboration Suite instance and its components, you can centrally manage all your Collaboration Suite instances through one tool (rather than through several Application Server Controls for Collaboration Suite) by using Grid Control. For example, suppose you have 10 Collaboration Suite instances deployed on 10 hosts. By deploying a Management Agent on each host, Enterprise Manager automatically discovers the key Collaboration Suite components on these hosts, and automatically begins monitoring them using default monitoring levels, notification rules, and so on.



---

---

## Identity Management

This chapter describes how you can use Grid Control to manage your Identity Management targets.

This chapter contains the following sections:

- [About Access and Identity Management](#)
- [Using Grid Control for Monitoring Identity Management Targets](#)
- [Identity Management Systems](#)
- [Identity Management Root Cause Analysis](#)
- [Automated Identity Management Monitoring and Alerts](#)
- [Diagnosing Identity Management Performance and Availability Problems](#)
- [Leveraging the Grid Control Management Framework](#)

### About Access and Identity Management

Access management is the means for controlling user access to enterprise resources. Access management products provide centralized, fine-grained access management for heterogeneous application environments, as well as out-of-the-box integration with Oracle products such as Oracle Portal, Oracle Collaboration Suite, and Oracle E-Business Suite.

Oracle Identity Management is a product set that allows enterprises to manage the end-to-end life cycle of user identities across all enterprise resources both within and beyond the firewall.

Automating user identity provisioning can reduce IT administration costs and improve security. Provisioning also plays an important role in regulatory compliance. Compliance initiatives focus on the enforcement of corporate policies as well as the demonstration of compliance with these standards. An enterprise identity management solution can provide a mechanism for implementing the user management aspects of a corporate policy, as well as a means to audit users and their access privileges.

Oracle Identity management products include the following:

- [Access Manager](#)
- [Identity Manager](#)
- [Identity Federation](#)

## Access Manager

Oracle Access Manager, formerly known as Oracle COREid Access and Identity, provides Web-based identity administration and access control to Web applications and resources running in a heterogeneous environments. It provides the user and group management, delegated administration, password management and self-service functions necessary to manage large user populations in complex, directory-centric environments. Access Manager supports all popular authentication methods including browser forms, digital certificates, and smart cards, and integrates seamlessly with most application servers and portals, including OracleAS 10g, BEA WebLogic, IBM WebSphere, Vignette and others. User identities and credentials can be accessed from a number of repositories including Oracle Internet Directory, Microsoft Active Directory, and Sun Java System Directory. With Access Manager, user access policies can be defined and enforced with a high degree of granularity through centralized management.

### Access Server

The Access Server enables you to protect resources such as URLs and legacy, non-HTTP applications. It provides authentication and authorization services to enterprise applications. It uses the information stored by the Identity Server to control which users, groups, and organizations can access a resource. It stores information about configuration settings and security policies that control access to resources in a directory server that uses Oracle Access Manager-specific object classes. You can use the same directory to store the Access Server configuration settings, access policy data, and user data, or you can store this data on separate directory servers.

### Identity Server

The Identity Server is a set of applications that provide delegated administration, user self-service, and real-time change management. The Identity Server stores information about users, groups, and organizations. For example, you can create, manage, and delete groups in the directory server. You can define a subscription policy for a group, including self-service with no approval needed, subscription with approvals, rule-based subscription, and no subscription allowed.

## Identity Manager

The Oracle Identity Manager platform automates user identity provisioning and deprovisioning and allows enterprises to manage the end-to-end lifecycle of user identities across all enterprise resources, both within and beyond the firewall. It provides an identity management platform that automates user provisioning, identity administration, and password management, wrapped in a comprehensive workflow engine.

Automating user identity provisioning can reduce IT administration costs and improve security. Provisioning also plays an important role in regulatory compliance. Key features of Oracle Identity Manager include password management, workflow and policy management, identity reconciliation, reporting and auditing, and extensibility through adapters.

Oracle Identity Manager also provides attestation support. Attestation is the process of having users or system managers confirm people's access rights on a periodic basis. Existing Sarbanes-Oxley requirements demand enterprises to perform attestation for all financially significant systems every three to six months. Identity Manager includes a highly flexible attestation solution to help enterprise customers meet these regulatory requirements in a cost-effective and timely manner. By setting up attestation processes in Identity Manager, enterprise customers can automate the

process of generation, delivery, review, sign-off, delegation, tracking, and archiving of user access rights reports for reviewers on a scheduled or ad-hoc basis.

## Identity Federation

As more companies move their business processes to the Web, many organizations have a greater need to extend the boundaries of their enterprise to include partner applications. Federated identity management allows companies to operate independently and cooperate for business purposes by enabling cross-domain single sign-on and allowing companies to manage user identities and vouch for them as they access resources managed by another domain.

Oracle Identity Federation, formerly known as COREid Federation, provides a self-contained federation solution that combines the ease of use and portability of a standalone application with a scalable, standards-based proven interoperable architecture. It helps corporations securely link their business partners into a corporate portal or extranet while also increasing their compliance with privacy and security regulations. Identity Federation enables companies to manage multiple partners and choose from industry standard federated protocols. Identity Federation provides built-in integration with customer's identity management infrastructure (Oracle and non-Oracle) to deliver an end-to-end user experience, addressing scenarios like automatic registration, identity mapping, seamless access control navigation, and others.

## Using Grid Control for Monitoring Identity Management Targets

Enterprise Manager helps you monitor the availability and diagnose the health of Access, Identity, Identity Federation, and Identity Manager servers within your enterprise configuration. By deploying a Management Agent on each host, you can use Enterprise Manager to discover the Identity Management components on these hosts, and automatically begin monitoring them using default monitoring levels, notification rules, and so on.

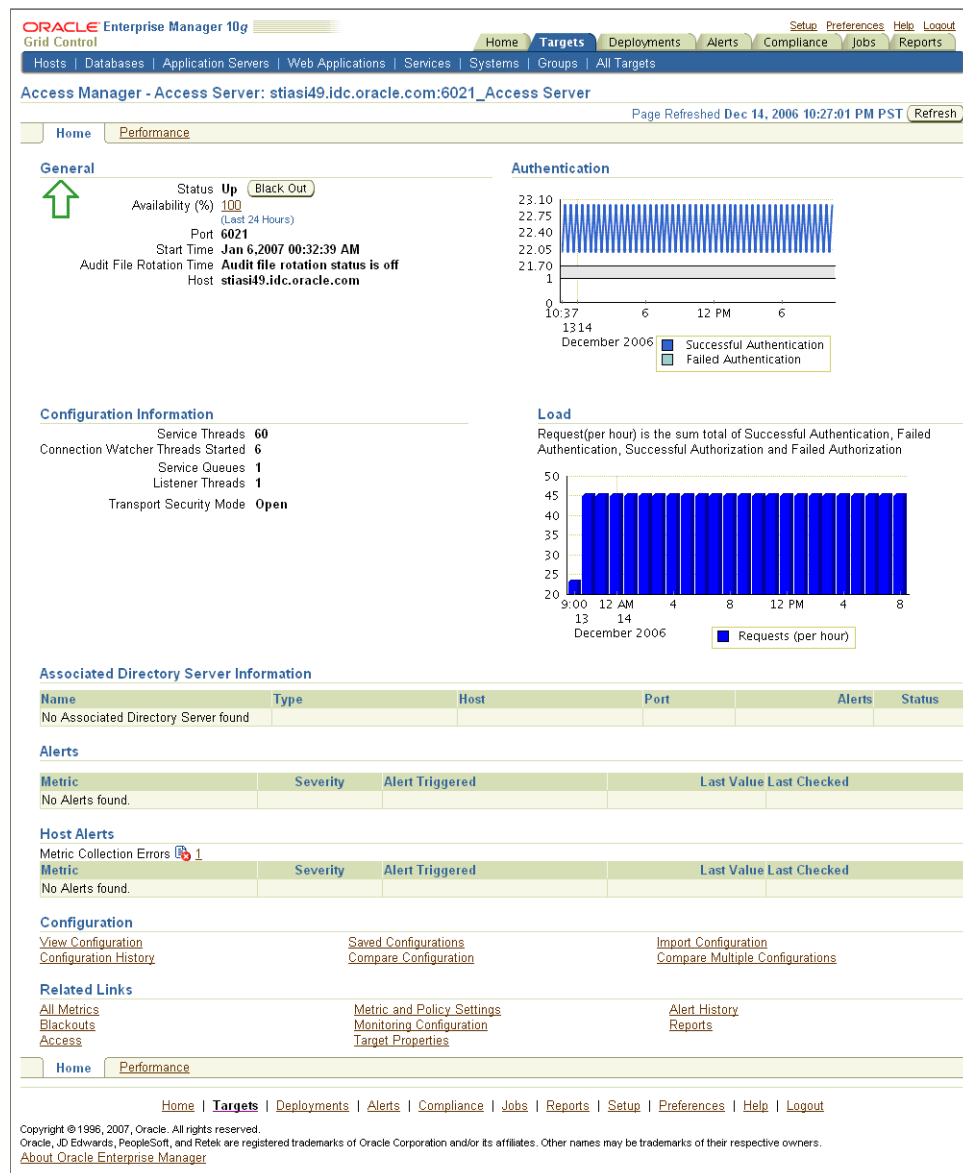
### Identity Component Server Home Page

All Identity Management targets, whether Access, Identity, Identity Federation, and Identity Manager have their own server home pages that provide easy access to key information required by the administrators. Each Identity Management Server home page provides the following information:

- Server status, responsiveness, and performance data
- Alerts and diagnostic drill-downs so you can identify and resolve problems quickly
- Resource usage for the server and its components
- Functionality to start, stop, and restart components in the case of locally-monitored Access and Identity servers
- Configuration parameters for Access and Identity servers

[Figure 11-1](#) shows the Access Manager - Access Server home page.

Figure 11–1 Access Manager - Access Server Home Page



## Identity Management Systems

Identity Management services run on Identity Management systems defined in Grid Control. The system includes the software infrastructure components that the Identity services rely on. The system includes components such as databases, HTTP servers, OC4Js, and other servers.

The system is a collection of server targets that are grouped together in Grid Control to give you a view of the "data-center" components that comprise your Identity Management deployment. Identity Management Systems are created when Identity suite components are discovered using Grid Control. Grid Control also monitors the performance and availability of these components and provides a System Dashboard to view the health of the Identity Management system in a single window.

Figure 11–2 shows an Access Manager-Identity System home page:

Figure 11–2 Identity System Home Page

ORACLE Enterprise Manager 10g  
Grid Control

Home | Targets | Deployments | Alerts | Compliance | Jobs | Reports


Hosts | Databases | Application Servers | Web Applications | Services | Systems | Groups | All Targets

Access Manager - Identity System: Identity System

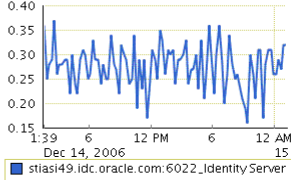
Page Refreshed Dec 15, 2006 1:24:50 AM Refresh Launch Dashboard

Home | Charts | Administration | Components | Topology

**Status** Up(1)



**ServiceTimePerRequest:Access Manager - Identity Server:Highest Average**



**Services**

Name	Type	Status	Performance Alerts	Usage Alerts	Policy Violations
(No data found.)					

**Configuration Changes**  
Configuration changes detected for the last 7 days

Category	Changes
All Target Types	0
Access Manager - Identity Server	0

**Load in the Last Collection Interval**

Request Processed 180  
Successful Cache Flush Request 0  
Cache Flush Request Failures 0  
Email Notification Failures 0  
Total LDAP Sessions opened 9

**Alerts**

Severity	Current	Last 24 hours
✘	0	0
⚠	0	0
Total	0	0

**Policy Violations**

Severity	Current	Last 24 Hours		Distinct Rules Violated
		Cleared	New	
(No violations found)				

**Security Policy Violations**

Severity	Current	Last 24 Hours		Distinct Rules Violated
		Cleared	New	
(No violations found)				

Security At a Glance

Critical Patch Advisories for Oracle Homes

Home | Charts | Administration | Components | Topology

**Related Links**

Access Edit System Properties  
Reports Target Properties

Home | Targets | Deployments | Alerts | Compliance | Jobs | Reports | Setup | Preferences | Help | Logout

Copyright © 1996, 2007, Oracle. All rights reserved.  
Oracle, JD Edwards, PeopleSoft, and Retek are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.  
About Oracle Enterprise Manager

## Configuring Identity Management Services

An Identity Management service is a logical target configured by Grid Control. You use Grid Control to step you through the process of configuring a web application service for your Identity component instances. After you configure a service, that service is displayed on the Services page.

Critical application functions are defined and monitored as services in Grid Control. Each service is monitored by Grid Control beacons, which run service tests that simulate real user access to the service. Service availability and performance are monitored automatically, and problems are immediately reported to the administrator. By monitoring availability and performance of Identity Management services, you can identify and resolve user-visible problems more quickly and thus minimize the impact on users.

### Creating an Access Service

Access service allows the Enterprise Manager administrator to perform service-level monitoring. If authentication or authorization services are not available, then the

administrator should be notified about the service failures. Administrators will then be able to diagnose the cause of the problem using Root Cause Analysis. For more information about Root Cause Analysis, see "[Identity Management Root Cause Analysis](#)".

1. Discover the Access Server target in Grid Control. For information about discovering this target, see the Grid Control online help. This will create the associated system.
2. In Grid Control, in the Services tab, select **Add Generic Service** and click **Go**.  
The Create Generic Service wizard is displayed.
3. In the Create Generic Service: General page, provide a name for the service and associate the Access system. Click **Next**.
4. In the Create Generic Service: Availability page, select **Define availability based on Service Test**. This allows you to create a web transaction. Click **Next**.
5. In the Create Generic Service: Service Test page, provide a name for the web transaction, select **Record a Transaction** and click **Go**.
6. In the Create Service Test page, in the Steps section, click **Record**.
7. In the Record Web Transaction page, click **Start**. A new browser window is displayed.
8. In the browser, enter the URL that is protected by the Access Server. In the Single Sign-On login page displayed, enter the username and password.
9. In the Record Web Transaction page, click **Stop** and then click **Continue**. The Create Web Applications: Service Test page is displayed again. Click **Continue**.
10. In the Create Generic Service: Beacons page, add a beacon to execute the web transaction you created. Click **Next**.
11. In the Create Generic Service: Performance Metrics page, click **Next**.
12. In the Create Generic Service: Usage Metrics page, click **Next**.
13. In the Create Generic Service: Review page, click **Finish** to create the Access service that will be monitored by the beacon test.

### Creating an Identity Service

Identity service allows the administrator to perform service-level monitoring of Access Manager Identity. If user management or group management services are not available, then the administrator should be notified about the service failures. Administrators will then be able to diagnose the cause of the problem using Root Cause Analysis. For more information about Root Cause Analysis, see "[Identity Management Root Cause Analysis](#)".

1. Discover the Identity Server target in Grid Control. For information about discovering this target, see the Grid Control online help. This will create the associated system.
2. In Grid Control, in the Services tab, select **Add Generic Service** and click **Go**.  
The Create Generic Service wizard is displayed.
3. In the Create Generic Service: General page, provide a name for the service and associate the Identity system. Click **Next**.
4. In the Create Generic Service: Availability page, select **Define availability based on Service Test**. This allows you to create a web transaction. Click **Next**.

5. In the Create Generic Service: Service Test page, provide a name for the web transaction, select **Record a Transaction** and click **Go**.
6. In the Create Service Test page, in the Steps section, click **Record**.
7. In the Record Web Transaction page, click **Start**. A new browser window is displayed.
8. In the browser, enter the URL of the Identity system (<host:port/identity/oblix>) and then click the Identity System Console link. The login page is displayed.
9. In the Record Web Transaction page, click **Stop** and then click **Continue**. The Create Web Applications: Service Test page is displayed again. Click **Continue**.
10. In the Create Generic Service: Beacons page, add a beacon to execute the web transaction you created. Click **Next**.
11. In the Create Generic Service: Performance Metrics page, click **Next**.
12. In the Create Generic Service: Usage Metrics page, click **Next**.
13. In the Create Generic Service: Review page, click **Finish** to create the Identity service that will be monitored by the beacon test.

### Creating an Identity Manager Service

For information about creating a web application service for Identity Manager, see the "Creating Web Applications for Identity Manager" topic in online help.

### Creating an Identity Federation Service

You can create an Identity Federation service for service-level monitoring of Identity Federation. The procedure to create an Identity Federation service is similar to the procedure to create an Access or Identity service.

## Monitoring Services

Grid Control enables you to monitor all of your Identity Management services. Each service is monitored for performance, usage, and availability.

Each service has its own home page. The Service Home pages in Grid Control provide:

- Status, responsiveness, and performance data
- Resource usage data for the service
- Summary information such as status, performance alerts, usage alerts, and policy violations for the service's subcomponents, including other services and associated systems
- Links to home pages for the service's subcomponents
- Alerts and diagnostic drill-downs so that you can identify and resolve problems quickly
- Services Dashboard

The Services Dashboard provides a high-level view of the status, performance, and usage of each Identity Management target. Service-level compliance for various time periods are also included for each service on the dashboard. You can launch the dashboard directly from Identity system target home page. You can also publish the Services Dashboard so that it can be viewed by non-Enterprise Manager users. This allows you to provide a self-service status web page to your end users.

- Related Links to do the following:
  - View metrics for the service
  - View client configurations
  - Edit the service
  - View the service target's properties
  - Manage blackouts
  - View and manage metric thresholds and policies

**See Also:** [Chapter 6, "Service Management"](#)

## Identity Management Root Cause Analysis

Individual services in Identity Management are associated with critical system components. This allows Enterprise Manager to perform Root Cause Analysis down to the system level whenever a service outage is detected. When you are configuring an Identity Management service in Grid Control, as mentioned in [Configuring Identity Management Services](#), you also mention the critical system components of this service. When an Identity Management service goes down, Enterprise Manager automatically performs a root cause analysis to determine which critical system component is responsible for this.

**See Also:** ["Diagnosing Service Problems"](#) in [Chapter 6, "Service Management"](#)

## Automated Identity Management Monitoring and Alerts

Enterprise Manager automatically gathers and evaluates diagnostic information from Identity Management targets distributed across the enterprise. As with all targets managed by Enterprise Manager, an extensive number of Identity Management performance metrics are automatically monitored against predefined thresholds. Alerts are generated in Grid Control when metrics exceed these thresholds.

## Diagnosing Identity Management Performance and Availability Problems

You can use Grid Control to diagnose performance and availability problems with your Identity Management services. For example, if a service outage occurs, Root Cause Analysis will determine if the primary cause is an outage of a critical service or system component. If a service performance issue is found, an administrator can examine detailed metrics over time related to that service and any of the service or system components used by that service. When you suspect there is a problem with one or more server components in the Identity Management system, the system home pages provide metrics and charts for diagnosing the issue.

**See Also:** ["Monitoring Services"](#) and ["Diagnosing Service Problems"](#) in [Chapter 6, "Service Management"](#)

## Leveraging the Grid Control Management Framework

Grid Control includes many general features that are useful to an Identity Management administrator, including:



- **Job Automation:** You can use the Grid Control job system to schedule tasks you want to automate.
- **Policies:** You can utilize the policy framework to ensure your Identity Management infrastructure adheres to your site-specific standards.
- **Database and Application Server Management:** Using the single Grid Control console, you can also manage the specific databases and application servers in your Identity Management deployment if needed.
- **Extensions:** Grid Control also includes monitoring of key network components that may be part of your Identity Management deployment. You can also extend Grid Control to monitor other components that are not recognized out-of-box by Enterprise Manager.



---

---

# Oracle Beehive Management

This chapter describes how you can use Grid Control to manage your Oracle Beehive targets.

This chapter contains the following sections:

- [About Oracle Beehive](#)
- [Using Grid Control for Monitoring Oracle Beehive Targets](#)
- [Discovering Oracle Beehive Sites](#)
- [Rediscovering an Oracle Beehive Site](#)
- [Oracle Beehive System and Services](#)
- [Automated Oracle Beehive Monitoring and Alerts](#)

## About Oracle Beehive

Oracle Beehive is the next generation of Oracle's enterprise collaboration technologies. Oracle Beehive provides a unified collaboration application and platform comprised of a comprehensive set of integrated, Java-based services, offering a new paradigm for enterprise collaboration solutions.

Oracle Beehive provides the full range of collaborative services for enterprise users, including:

- E-mail
- Time management
- Instant messaging
- Content and document management
- Task management
- Discussion forums
- Search
- Voicemail
- Contact management
- Mobile device support
- Online presence

The following terms and concepts form the foundation of the Oracle Beehive system, enabling many of the collaboration-based features that it provides:

- Oracle Beehive Service and Service Instances – Oracle Beehive offers discrete, function-specific services that interoperate seamlessly to provide a wide range of features for enterprise collaboration, compliance, third-party component connectivity, coexistence, and so on. In Oracle Beehive, the term service instance refers to a single instantiation of a service on a particular server. The term service typically refers to the collection of all instances of a particular service across all of the servers that host them. In Enterprise Manager, an Oracle Beehive service is called an Oracle Beehive Application and an Oracle Beehive service instance is called an Oracle Beehive Application instance.
- Site – A site is a collection of physical hardware in the same geographic location used to run Oracle Beehive. Each Oracle Beehive site requires a minimum set of components to support a fully-functioning system. However, within each site, multiple instances of a variety of supported components, including services and servers, can be implemented.
- Workspace – A workspace is both a virtual location and a container that provides Oracle Beehive users a place to collaborate, and enables them to create, view, store, access, and manage the artifacts related to their collaborations. From the user perspective, workspaces appear in supported Oracle Beehive clients.

## Using Grid Control for Monitoring Oracle Beehive Targets

Enterprise Manager helps you monitor Oracle Beehive components in your enterprise. You can discover Oracle Beehive components in your environment and add them to Grid Control for central monitoring and management. Enterprise Manager Agents will discover and monitor Oracle Beehive sites, services, and service instances. In Grid Control Oracle Beehive sites, services (called applications), and service instances (called application instances) have home pages.

Following are the target types defined for Oracle Beehive in Grid Control:

- Site – A site is a top-level container for Oracle Beehive deployment in a particular geographic location.
- CalDAV Application and CalDAV Application Instance – Oracle Beehive CalDAV Service supports and manages all time management-related features and settings that the system leverages over the Calendaring Extensions to WebDAV (CalDAV) protocol. Enterprise Manager collects monitoring metrics from all Oracle Beehive CalDAV service instances. These represent metrics for CalDAV application instances. The metrics from instances are aggregated for CalDAV application metrics. Enterprise Manager also collects Configuration parameters for CalDAV application.
- Conference Application and Conference Application Instance – Oracle Beehive Conference Service manages all aspects of Web-based meetings and conferences for Oracle Beehive, enabling meeting organizers and participants to conduct collaborative sessions online through Oracle Beehive workspaces. Enterprise Manager collects monitoring metrics from all Oracle Beehive Conference service instances. These represent metrics for Conference application instances. The metrics from instances are aggregated for Conference application metrics. Enterprise Manager also collects Configuration parameters for Conference application.
- Device Management Application and Device Management Application Instance – Oracle Beehive Device Management application enables system administrators to manage supported client software installed on computers, such as Oracle Beehive Integration for Outlook, and mobile devices. Enterprise Manager collects

monitoring metrics from all Oracle Beehive Device Management service instances. These represent metrics for Device Management application instances. The metrics from instances are aggregated for Device Management application metrics. Enterprise Manager also collects Configuration parameters for Device Management application.

- Discussions Application and Discussions Application Instance – Oracle Beehive Discussions application enables organizations to host threaded, online discussion forums in which users can browse message boards, and post and respond to messages. Enterprise Manager collects monitoring metrics from all Oracle Beehive Discussions service instances. These represent metrics for Discussions application instances. The metrics from instances are aggregated for Discussions application metrics.
- E-mail Application and E-mail Application Instance – Oracle Beehive E-mail application supports all aspects of e-mail creation, delivery, and management for Oracle Beehive, by exposing e-mail from coexisting systems to supported clients. The E-mail application enables users to organize and manage their e-mail messages in the context of Oracle Beehive workspaces. Enterprise Manager collects monitoring metrics from all Oracle Beehive E-mail service instances. These represent metrics for E-mail application instances. The metrics from instances are aggregated for E-mail application metrics. Enterprise Manager also collects Configuration parameters for E-mail application.
- FTP Application and FTP Application Instance – Oracle Beehive FTP application supports and manages all content management-related features and settings that the system leverages over the File Transfer Protocol (FTP) and the Secure File Transfer Protocol (FTPS). Enterprise Manager collects monitoring metrics from all Oracle Beehive FTP service instances. These represent metrics for FTP application instances. The metrics from instances are aggregated for FTP application metrics. Enterprise Manager also collects Configuration parameters for FTP application.
- IM (Instance Messaging) Application and IM Application Instance – Oracle Beehive Instant Messaging (IM) application supports all aspects of instant messaging for Oracle Beehive. With the Instant Message application, users can send text-based messages in real time. Users can also attach files to messages, and can broadcast messages to multiple recipients simultaneously, even to entire groups.
- Message Delivery Application and Message Delivery Application Instance – Oracle Beehive Message Delivery application handles all aspects of routing and channel support for Oracle Beehive notifications through Short Message Peer-to-Peer (SMPP), SMTP, and SMS. The Message Delivery application also provides built-in e-mail, instant messaging, SMS, and voice message delivery handlers that can be enabled or disabled, as needed.
- Presence Application and Presence Application Instance – Oracle Beehive Presence application supports and manages all aspects of user and resource presence for Oracle Beehive. Presence is the ability to detect and identify the status of a user or resource, and then display that status to other users and resources.
- PushMail Application and PushMail Application Instance – Oracle Beehive Mobile Mail application manages the features and settings related to push mail for supported mobile devices. The Mobile Push application enables communication between device-resident, device management clients, and the Oracle Beehive Device Management application. The Mobile Push application also manages the features and settings that are related to the delivery of notifications to supported mobile devices.

- Search Application and Search Application Instance – Oracle Beehive Search Service supports and manages all aspects of user-initiated, text-based searches for Oracle Beehive. It enables users to search for e-mail messages, documents, and meeting entries. Enterprise Manager collects monitoring metrics from all Oracle Beehive Search service instances. These represent metrics for Search application instances. The metrics from instances are aggregated for Search application metrics. Enterprise Manager also collects Configuration parameters for Search application.
- Time Management Application and Time Management Application Instance – Oracle Beehive Time Management application supports all aspects of time and task management, as well as user and resource scheduling for Oracle Beehive. The Time Management application maintains and provides access to the latest schedule information for all Oracle Beehive users and resources.
- WebDAV Application and WebDAV Application Instance – Oracle Beehive WebDAV application supports and manages all content management-related features and settings that the system leverages over the Web-based Distributed Authoring and Versioning (WebDAV) protocol.
- Voice Message Application and Voice Message Application Instance – Oracle Beehive Voicemail application supports all aspects of voicemail and fax management for Oracle Beehive. The Voicemail Service can be leveraged by the E-mail Service, enabling delivery of voicemail messages and faxes as e-mail.
- Workspaces Application and Workspaces Application Instance – Oracle Beehive Workspace application supports all the features and functionality provided by Oracle Beehive personal and team workspaces. Workspaces are the core of the user experience with Oracle Beehive, especially in regard to the collaborative activities of teams. Therefore, the Workspace Service is responsible for consolidating and exposing, in a single location, the collaborative functionality provided by the other Oracle Beehive applications, including the E-mail application, the Time Management application, the Instant Message application, the Discussions application, the Meetings application, the Voicemail application, the Search application, and so on.
- User Directory Application and User Directory Application Instance – Oracle Beehive User Directory application supports all aspects of user directory management for Oracle Beehive, including integration with existing Oracle and third-party user directories.
- XMPP Application and XMPP Application Instance – Oracle Beehive XMPP application supports all aspects of XMPP for Oracle Beehive including support for XMPP 0.9 and 1.0.

## Discovering Oracle Beehive Sites

Enterprise Manager can discover Oracle Beehive sites from remote as well as local agents. Local agents could be used to monitor application instances. Oracle Beehive site, application and application instances discovery, monitoring and configuration parameter collections use the JMX interface exposed by Oracle Beehive's BEEMGMT OC4J. All Oracle Beehive application instances are grouped under the corresponding Oracle Beehive application type. For example, all E-mail application instances will be grouped under a single E-mail application.

To add new Oracle Beehive targets to be monitored by Grid Control, you will need to initiate the discovery process for each Oracle Beehive site.

## Oracle Beehive Site Home Page

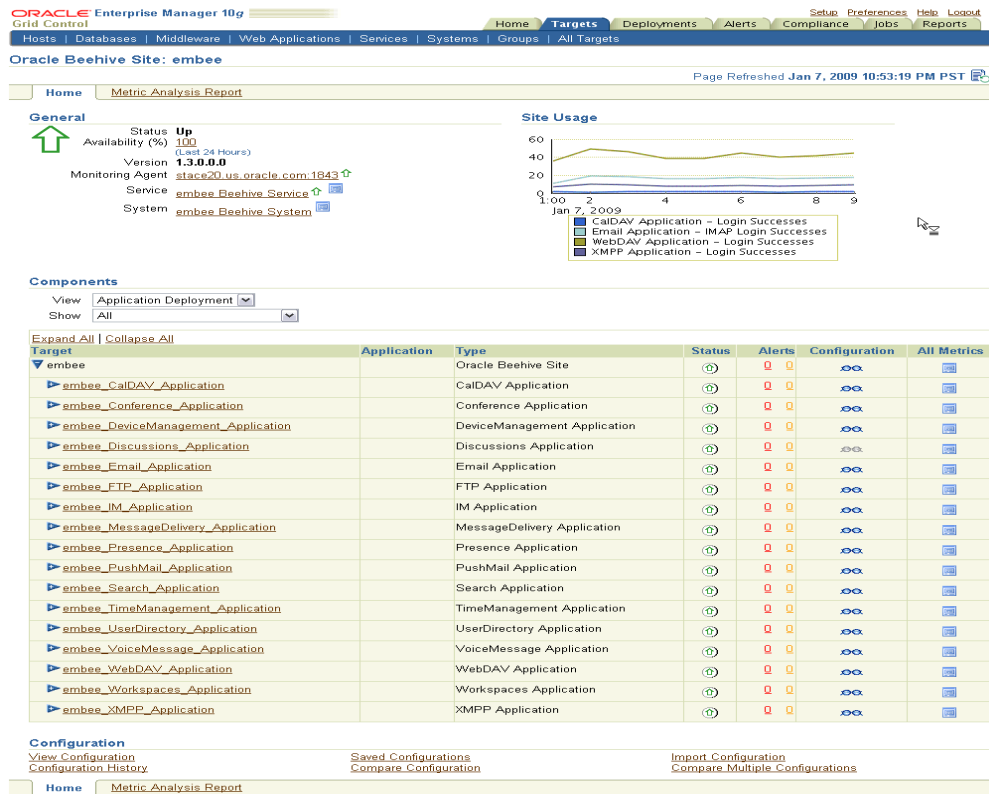
The Oracle Beehive Site Home page provides information about:

- All Oracle Beehive applications and application instances within the Oracle Beehive site
- Current Status, Availability, and Site Version
- Site usage for commonly used Oracle Beehive Services including CalDAV, WebDAV, E-mail (IMAP), and XMPP
- Configuring Systems and Services
- Accessing System and Services dashboards for snapshots of the system and service-related information
- Drill-downs to view monitoring metrics and configuration for each Oracle Beehive application and application instance
- Alerts and diagnostic drill-downs so you can identify and resolve problems quickly

You can filter the view based on Physical or Application Deployment of Oracle Beehive. Application Deployment refers to the Oracle Beehive services deployed in an Oracle Beehive site. There can be multiple instances of Oracle Beehive services deployed within one or more OC4J containers. You can further filter the Application Deployment view by applications. Physical Deployment view shows the hosts, Oracle Application Servers, HTTP servers, OC4Js, and the database comprising an Oracle Beehive site. You can further filter this view based on the host, database, or platform components.

[Figure 12-1](#) shows the Oracle Beehive Site home page.

Figure 12–1 Oracle Beehive Site Home Page



If the Oracle Beehive System and Services have not been created, you can create them by configuring the service tests for the Oracle Beehive Site.

Once the Oracle Beehive System and Services is configured, you can view snapshots of the service and system related information from the service and system dashboards on the Oracle Beehive Site home page. For more information about Oracle Beehive System and Services, see "Oracle Beehive System and Services". For information about Systems and Services in Enterprise Manager, see Chapter 6, "Service Management".

## Rediscovering an Oracle Beehive Site

Enterprise Manager allows you to rediscover an Oracle Beehive site and its sub-components. During rediscovery, all sub-components of Oracle Beehive are compared. Rediscovery of Oracle Beehive includes addition of new Oracle Beehive targets, deletion of Oracle Beehive targets, or changing the configuration of existing Oracle Beehive targets.

**Note:** Ensure that you have the necessary privileges to create, delete, and view targets.

## Oracle Beehive System and Services

Since Oracle Beehive offers a set of complex and critical applications and services, it is very important for IT organizations to monitor and manage them for availability and performance. Enterprise Manager Oracle Beehive Site Systems and Services help you to effectively manage Oracle Beehive services from the overview level to the individual component level. Using Grid Control, you can identify and map service



unavailability and performance degradation to one or more component failures. Grid Control not only monitors individual components in Oracle Beehive deployments such as hosts, databases, OC4Js, application servers, Beehive Site, Beehive application instances (for example, E-mail application instances, Instant Messaging application instances), etc., but also the applications hosted by those components, allowing you to monitor critical applications using the top down approach and from end-user perspective.

A "Service" in Enterprise Manager is defined as an entity that provides a useful function to its users. It models a business process or application, such as e-mail, time management, instant messaging, and so on. A "System" is a group of underlying components, such as hosts, databases, OC4Js, application servers, etc., on which the Service runs. A "Service test" is a functional test that is run against a service to determine the availability and performance of a service. For example, an IMAP Service test would determine the availability and performance of the E-mail service accessed via standard IMAP protocol. Availability of a service is a measure of the end-users' ability to access the service at a given point in time. Performance indicates the response time as experienced by the end-users. For information about Services and Systems, see the Service Management and System Management chapters in this guide.

Following is the complete Services hierarchy included in the model template. The hierarchy shown in the wizard is subject to the Oracle Beehive applications deployment. For example, if there is no FTP application discovered during Oracle Beehive Site discovery, the wizard will not include FTP Service in the hierarchy.

- Beehive Site Service
  - Email Service
    - \* IMAP Service
    - \* SMTP Service
  - Time Management Service
    - \* CalDAV Service
  - Presence and IM Service
    - \* XMPP Service
  - Conferencing Service
  - Mobile Service
    - \* Mobile Mail Service
  - Client Service
    - \* Zimbra Service
  - WebDAV Service
  - FTP Service
  - Voice Message Service
  - Workspaces Service
  - User Directory Service

Following is the complete Systems hierarchy included in the model template:

- Beehive Site System
  - Core System

- Email System
- Time Management System
- Presence and IM System
- Conferencing System
- Mobile System
- WebDAV System
- FTP System
- Voice Message System
- Workspaces System
- User Directory System

Based on the System and Services hierarchy configured through the wizard, an Oracle Beehive Site Service dashboard and an Oracle Beehive Site System dashboard will be automatically generated.

---

---

**Note:** The Core System will be created only if there are agents monitoring the Oracle Beehive hosts, Oracle application servers, OC4Js, databases, LDAP servers, and so on.

---

---

For information about Oracle Beehive service dashboard metrics, see [Appendix A](#).

## Automated Oracle Beehive Monitoring and Alerts

Enterprise Manager automatically gathers and evaluates diagnostic information from Oracle Beehive targets distributed across the enterprise. As with all targets managed by Enterprise Manager, an extensive number of Oracle Beehive performance metrics are automatically monitored against predefined thresholds. Alerts are generated in Grid Control when metrics exceed these thresholds.

---

---

## Lifecycle Management

This chapter gives an overview of the server and software provisioning and patching features introduced in Enterprise Manager 10gR2. This chapter contains the following:

- [Lifecycle Management Overview](#)
- [Software Provisioning](#)
- [Server Provisioning](#)
- [Provisioning Concepts](#)
- [Using the Provisioning Features](#)
- [Patching](#)

Provisioning and patching features of the Enterprise Manager together make up the Lifecycle Management (Grid Automation) solution area of the Enterprise Manager. To read more about this solution area visit:

[http://www.oracle.com/technology/products/oem/mgmt\\_solutions/provisioning.html](http://www.oracle.com/technology/products/oem/mgmt_solutions/provisioning.html)

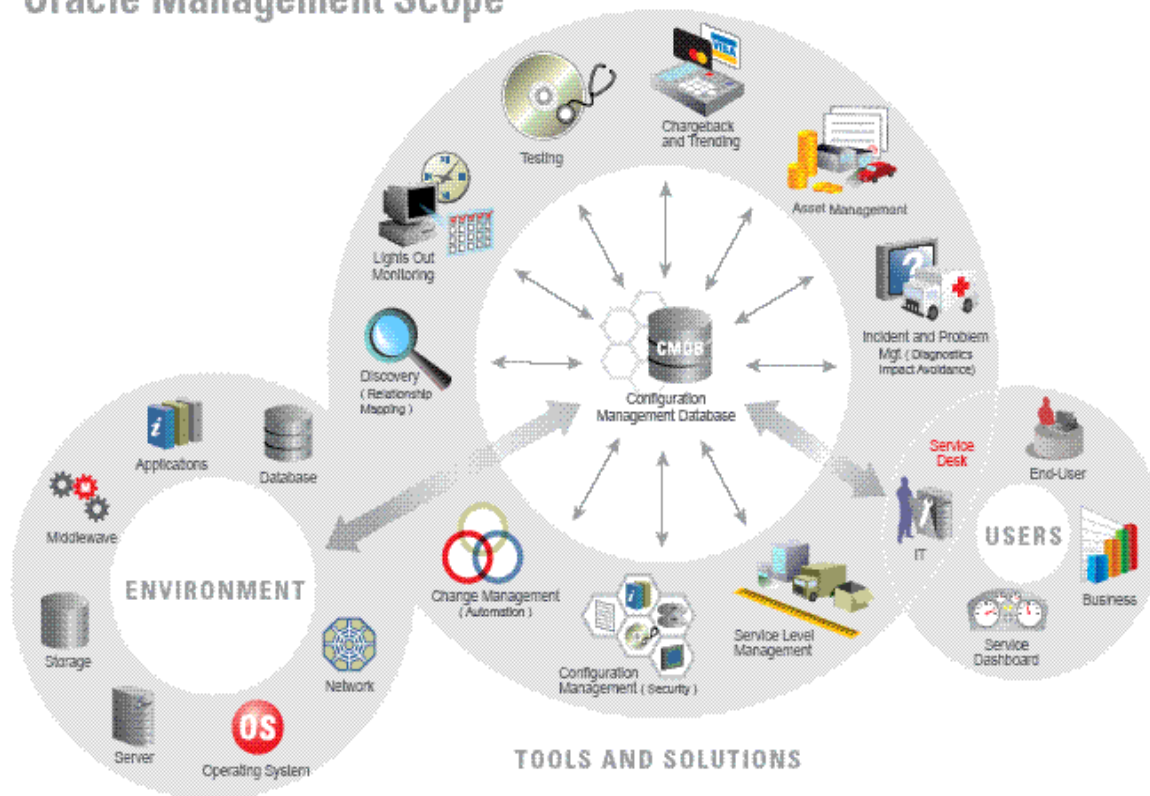
### Lifecycle Management Overview

The provisioning and patching features of Enterprise Manager automate the deployment of software, applications, and patches. They make critical data center operations easy, efficient, and scalable resulting in lower operational risk and cost of ownership. The ability to provision and patch the entire software stack that includes the operating system, the middleware, database, third party software, and applications supplemented by comprehensive reporting tools make these features extremely significant entities in the overall System Management space.

As shown in [Figure 13-1](#), Enterprise Manager covers the entire lifecycle management of software, applications, and servers. Enterprise Manager orchestrates the initial reference sandbox deployment and then the mass unattended deployment of gold images created from these reference deployments. The smaller lifecycle shown transcribed in the figure automates the ongoing patch lifecycle management of the various deployments. Right from proactively informing the administrators about the critical patches and vulnerabilities in the deployments, acquiring these patches to mass deployment and verification of these patches is automated by the Enterprise Manager. Going forward, as the computation demand for the resources decline, Enterprise Manager allows one to deactivate and de-provision the resources making them available for a different purpose.

Figure 13-1 Lifecycle Management Overview

## Oracle Management Scope



### Advantages of Using Enterprise Manager Provisioning and Patching Features

Following are the advantages of using the provisioning and patching features in Enterprise Manager:

- The provisioning and patching features provide a repeatable, reliable, and automated solution for performing mass, unattended, and schedulable deployment of
  - Software and servers based on Gold Images created using reference deployment or installation media
  - Software and operating system updates
  - Complex and multi-tier software like Oracle Real Application Clusters (RAC) and Application Server Clusters
- Orchestrates not only provisioning of software but completely automates configuration of software and ensures zero-time for patching of mission critical

systems by orchestrating rolling patching for complex multi-tier installation like Real Application Clusters (RAC) databases and Application Server clusters.

- Allows new resources to be provisioned at short notice based on compliant and tested gold images.
- Enterprise Manager's Critical Patch Facility proactively and regularly queries My Oracle Support for critical patches that have been released and notifies the administrators of only those patches applicable to them. It can also invoke the Patch application in context, and remedy the vulnerable installations. The Critical Patch Facility also supports an offline mode to serve the case of data centers that are not connected to the Internet.
- Completely automates the patching operations across the stack. For example, for database patching, it takes care of shutting down and starting up database instances as required by the patch.
- Allows multiple operations to be accommodated in a single change window.
- Enterprise Manager also provides command-line interface support to all out-of-box provisioning and patching deployment procedures. These features can hence be invoked by custom scripts.
- Software provisioning and patching features support SUDO and PAM authentication.
- Single interface for multiple players. For example, component designers responsible for creating gold images based on corporate standards and the operators all use the same Enterprise Manager console.
- Automation of repeatable installation and patching operations across the stack leads to substantial cost savings in terms of costs and man-hours.

## Wide Coverage Across the Stack

Enterprise Manager provides provisioning and patching capabilities across the stack for:

- Operating Systems, with Bare Metal Provisioning on Linux and operating system patching
- Databases, with Real Application Clusters (RAC) provisioning, extension, and deletion and flexible patching for Oracle database and Oracle Real Application Clusters
- Middleware, with application server J2EE, BPEL, and SOA provisioning and patching

Once the above entities are available in the software library, one of the following features can be used mass deployment. Note that these features require Oracle Management Agents to be present on the target machines where the software has to be provisioned.

## Software Provisioning

Enterprise Manager provides the feature of provisioning software or applications.

## Deployment Procedures

The workflow of all the tasks that need to be performed for a particular life cycle management activity is encapsulated in a Deployment Procedure. A Deployment

Procedure is a hierarchal sequence of provisioning steps, where each step may contain a sequence of other steps. It provides a framework where specific applications and procedures can be built. Oracle ships a set of best practices Deployment Procedures to accomplish provisioning and patching related tasks. Deployment Procedures can be extended and customized for customer needs. The Deployment Procedure to patch a single instance database differs from the one to patch a RAC environment or an Application Server.

Deployment Procedures can vary from one customer to another or a test installation to a production installation. Deployment Procedures take into account and resolve the reality that environments are often different with each having complexities across different tiers with multiple dependencies. The situation is further compounded by existing operational practices. In a typical data center, Deployment Procedures can involve a design time activity (typically performed by a Lead Administrator) and a runtime activity (typically performed by the Operator).

For more information about deployment procedures, see the Deployment Procedures Best Practices White Paper and "Using Enterprise Manager For Grid Automation With Deployment Procedures" in the Oracle Enterprise Manager Advanced Configuration Guide available on OTN.

## Cloning

The cloning application was used during the earlier versions of Enterprise Manager to clone Oracle software either from reference installation or using gold images in software library. These have been replaced by the more flexible deployment procedures.

Use the cloning feature to clone older versions of applications such as Application Server 9.2.0.6. Use the cloning wizard available from the Grid Control Console to perform cloning operations.

It is recommended that you use deployment procedures for provisioning and patching operations.

For more information about cloning, see Cloning, Provisioning RAC & AS Environments Using Enterprise Manager 10g based Cloning.

---

---

**Note:** For cloning, the Cloning wizard will not be supported from Enterprise Manager version 10.2.0.4.

---

---

## Server Provisioning

Bare metal or Operating System provisioning application provides server lifecycle management to build, manage, and optimize server infrastructure. The application:

- Automates deployment of consistent, certified Linux operating system images along with larger number of servers.
- Leads to faster, unattended deployment of software and operating systems.
- Allows provisioning of middleware, Clusterware, Real Application cluster (RAC) etc.
- Provides a template-based approach for provisioning a variety of Linux configurations servers (RedHat 3.0/4.0, SuSE/SLES9). This also ensures compliance to standards and consistency across all deployments.
- Reduces errors with standardized gold image-based server provisioning.

- Supports heterogeneous hardware and network configuration.
- Automatically discovers bare metal and live target servers for provisioning.
- Especially for Oracle software, the application encodes best practices out-of-the-box for patching.
- Results in considerable reduction in manual labor that leads to substantial cost savings.

---

---

**Note:** Ensure that the targets being provisioned are managed by agent versions 10.2.0.2 or higher.

---

---

For detailed use cases and capabilities of the Bare Metal Provisioning application, refer to the Best Practices for Grid Control based Bare Metal Provisioning White Paper.

The above-mentioned provisioning deployment procedures can also be executed through EMCLI. For more information, see Chapter 2 in the Oracle Enterprise Manager Advanced Configuration Guide.

## Provisioning Concepts

Following are the basic elements associated with provisioning.

### Components

Components represent the primary building blocks that may be combined with other components as needed, to specify the complete software configuration or image that is provisioned on target machines. A component can represent operating system software, Oracle software or any third party software and applications. Software components are individually maintained within the Oracle Software Library. Versions, states, and maturity levels can be associated with each component.

### Directives

Directives can be imagined as instructions to cook the final image (recipe) using components (ingredients). These are constructs used to associate scripts with software components and images. These scripts contain directions on how to interpret and process the contents of a particular component or an image. Directives encapsulate the script, the command line used to invoke the script, and the script configuration properties. They capture everything required for invoking the script on a machine during a provisioning operation. Directives are usually categorized based on the provisioning life cycle phases they are targeted for, or the actions they perform. Imagine Directives as set of executable instructions that run from a supported shell (for example, *borne-again*, Perl, Python), programming language (for example, Java), or execution framework or interpreter (such as “*make*” or “*ant*”). Directives are contained within a file stored in the Oracle Software Library and referenced from the software components that employ them.

Components and Directives are used by Deployment Procedures (both out-of-box and custom procedures) to mass deploy software and applications on to target servers.

### Images

An image can be viewed as a set of components and may include directives that form the required software configuration, which is deployed on the target machines. An

image contains the complete software stack from operation system to application, in the form of its components and is used for provisioning servers from ground up with the entire stack provisioned on them. Images reference the components they logically contain by version (rather than include them directly). Images are stored in the Oracle Software Library and versions, states, and maturity levels can be associated with them.

## Software Library

Enterprise Manager allows a shared location accessible from the Oracle Management Server (OMS) to serve as a Software Library. Software library serves as the central repository for metadata and binary content for components, images, and directives. It allows maintaining versions, maturity levels, and states of components, directives, and images.

---

---

**Note:** For server provisioning, other basic elements like Network Profiles and Assignments are required. Refer to the Concepts section in the Best Practices for Grid Control based Bare Metal Provisioning White Paper.

---

---

## Using the Provisioning Features

Following are the one-time configuration activities for using the provisioning features.

### Setting up and configuring the environment

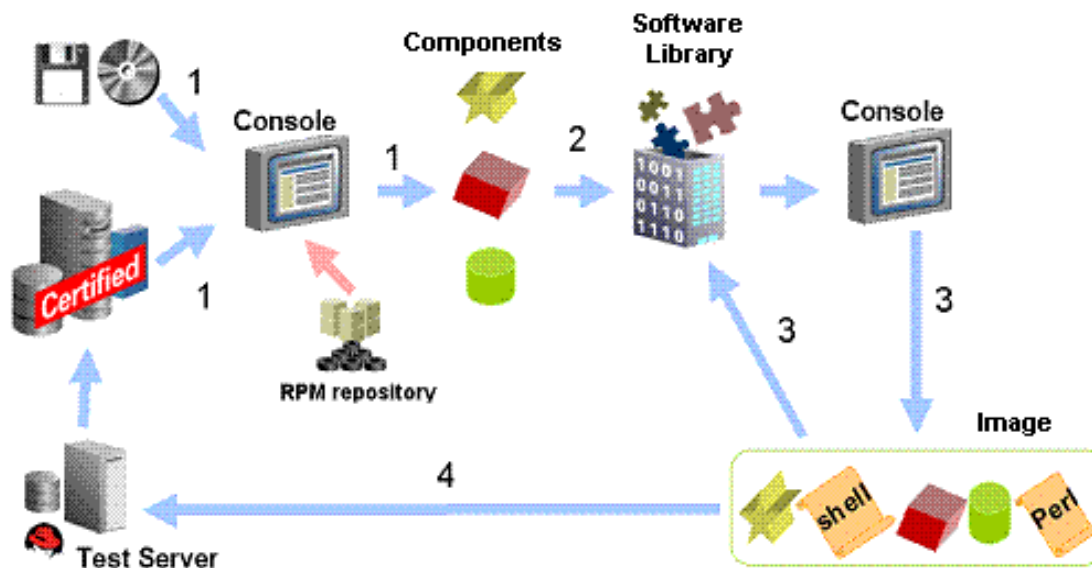
For both software and server provisioning the user needs to perform a one-time activity of setting up a Software Library. For server provisioning additional elements like Boot server, Stage server, and RPM repository have to be configured as required by the provisioning application. Once configured, the same elements will be used for any software or server-provisioning operation performed using the provisioning application.

### Creating Components, Directives, and Images

Once the environment is ready, the user can use the Enterprise Manager user interface to create components, directives, or images for deploying them onto the target servers. This is explained in [Figure 13-2](#).



**Figure 13-2 Creating Software Components Using Enterprise Manager User Interface**



- You can use either tested reference installations or installation media to create software components from the Enterprise Manager User Interface. The RPM repository is used for creating the out-of-box operating system components that one needs to provision on the bare metal or live servers. You can use the Enterprise Manager User Interface to create Directives and other server provisioning constructs like Storage templates, Hardware templates, and Network templates.
- The reusable entities created above are stored in the Software Library.
- These reusable entities can then be used by deployment procedures for deployment or mixed and matched to create deployable images for the hardware servers, which are again stored in the software library.
- The images or components can then be deployed on test or production environments.

## Patching

Manually applying software patches to maintain the latest and most secure IT environment can become a full-time job. With Enterprise Manager's deployment management tools, you can quickly see the patches available for the components in your enterprise, find out which have not been applied and which are critical, then bring those deployments up to the latest patch level with out-of-box best practices.

The enriched patching application offers an "end-to-end" patching solution that works seamlessly across a wide range of product patches and customer environments. The patching application automates the deployment of patches for the Oracle Database, including Clusterware and Oracle RAC, as well as Oracle Application Server. Also, Out-of-box procedures are provided for patching Operating Systems - Linux (Oracle Enterprise Linux, RHAT, and SUSE), Solaris, and Windows.

Using a direct link to My Oracle Support patch repository, the Critical Patch Facility identifies the critical patches that have been released for the Oracle software running in your specific systems, and notifies administrators of only those patches that are

applicable to their environment. Once a patch is identified, Grid Control can orchestrate the download and deploy it on multiple targets automatically.

Enterprise Manager provides the following patching features:

- [Patching Through Deployment Procedures](#)
- [Patching Features](#)
- [Application of Critical Patch Updates](#)
- [Application Server Patching](#)
- [Linux Host Patching](#)
- [Operating Systems Patching](#)
- [Accessing Patching Pages in Grid Control](#)

## Patching Through Deployment Procedures

Deployment procedures are the best practices for orchestration of patching Oracle Software such as Databases including Real Application Clusters, Clusterware, Automated Storage Management, Application Servers, and Operating Systems. The Deployment procedure-based infrastructure has been leveraged to increase the power and flexibility of Oracle patching for complex multi-tier environments. The out-of-box Deployment Procedures are Oracle-provided best practices that can be customized for specific needs. Users can enable and disable or add custom steps for specific actions and create the best practice for their environment. This activity is a one-time design activity typically of the lead DBA, which can be the standard and carried over by the operator for the entire environment.

Deployment procedures also support secure host authentication using sudo or PAM. The entire exercise can be run in the command line (CLI) mode, thereby making it possible to integrate with the existing scripts. Refer to Enterprise Manager Advanced Configuration for details on Using Deployment Procedures.

---



---

**Note:** For patching Oracle Management Agents, use the patch wizard, which can be accessed by clicking the **Patch Agent** link under Patching section in the Deployments page.

---



---

## Patching Features

Grid Control's patching tools simplify the patching of Oracle software products. Some key features are listed in [Table 13–1](#).

**Table 13–1 Features of Patching**

Feature	Description
Software Library	The Library to store patches, directives or components. Can be used in offline mode of patching. You can upload patches to Software library using the View/Upload Patch link.  Refer to <i>Using the Software Library</i> section in <i>Oracle Enterprise Manager Advanced Configuration</i> for details on Software Library.

**Table 13–1 (Cont.) Features of Patching**

<b>Feature</b>	<b>Description</b>
Critical Patch Advisories	<p>Lists all critical advisories with their corresponding areas of impact.</p> <p>Critical Patch Advisories also provides support for "remedies," in that you can select an advisory and view the calculated remediation paths from the context of that advisory, as well as the affected Oracle homes.</p>
Online Patching	Allows you to connect to My Oracle Support through Grid Control, search and download the required patches, and apply.
Offline Patching	Allows you to perform all the patching activities through Software Library. Even when you are not connected to My Oracle Support, you can search, download, and apply patches.
My Oracle Support	Searches My Oracle Support Web site for Oracle patches and patch sets. Or use Grid Control to search after you provide your My Oracle Support Web site user name and password.
Patching Setup	<p>Helps you configure My Oracle Support, patching, proxy connection, and offline patching settings.</p> <p>Note that if you are accessing a proxy server to get to My Oracle Support, you will need to provide proper authentication and credentials.</p>
Patching through Deployment Procedure	<p>Oracle ships a set of best practices Deployment Procedures to accomplish provisioning and patching related tasks. Deployment Procedures can be extended and customized for customer needs. This allows:</p> <ul style="list-style-type: none"> <li>▪ SIDB, RAC/Clusterware patching in rolling modes</li> <li>▪ Helps you stage patches (and also to the central location) and apply them to the required Oracle homes</li> <li>▪ When patching database by targets, provides additional options that allow you to shut down the database, apply SQL scripts (and start the database in Upgrade mode), and start up the database again</li> <li>▪ Integration with Notification systems for various statuses</li> </ul>
Agent Patching	Patching Oracle Management Agents through the Patch Agent link. This applies agent-specific patches and also generic patches like CORE or DST patches on the Agent.
Shared Agent Patching	Automates patch applications on shared agents when they are NFS-mounted. During Shared Agent Patching, patches the central location where the agent is installed, shuts down and starts up the shared agents, and executes any pre/post-patching scripts (if specified).
Support for SUDO/PAM-based patching	Deployment procedures Secure host authentication for patching using SUDO/PAM.
Notification	This feature notifies users by identifying Criticality on the targets across the Target pages.
Oracle Home Credentials	When you override preferred credentials, you can choose to either specify one set of credentials for all Oracle homes, or specify different credentials for each home.

**Table 13–1 (Cont.) Features of Patching**

Feature	Description
Information Publisher	<p>Provides a powerful central reporting framework that produces detail and summary reports on patch deployments and non-compliant installations. Supports both out-of-box and ad hoc reporting to satisfy different customer needs.</p> <p>The Library to store patches, directives or components. Can be used in offline mode of patching. You can upload patches to Software library using the View/Upload Patch link.</p> <p>Refer to <i>Using the Software Library</i> section in the <i>Oracle Enterprise Manager Advanced Configuration</i> for details on Software Library.</p>

## Application of Critical Patch Updates

You can use Grid Control to manage Oracle Critical Patch updates:

- **Assess Vulnerabilities:** This helps identify the Oracle Software affected by the advisory. The list displays comprehensive detail on the Critical Patch applicable on specific products under version and platform. This also displays the affected Oracle Homes.

Grid control automates the entire process of critical patch application. It performs an assessment of vulnerabilities by examining your enterprise configuration to determine which Oracle homes have not applied one or more of these critical patches. Grid Control provides a list of critical patch advisories and the Oracle homes to which the critical patches should be applied.

Some Oracle software patches have been identified as critical. To help ensure a secure and reliable configuration, all relevant and current critical patches should be applied to the appropriate targets in your enterprise.

From the summary of patch advisories, you can navigate for more information about a particular patch, and get a list of the Oracle homes to which the patch has not been applied. Then you can launch the Grid Control Patch tool to download and deploy the patches to multiple targets.

- **User Notification:** This feature notifies users by identifying Criticality on the targets across the Target pages. Also, with the notification and reporting one can receive notifications and reports for the Critical Patch advisory and its assessment.
- **Application of critical patches:** The patch application process is automated directly from the assessment and the patch is downloaded from My Oracle Support and orchestrated:
  - Directly automate the patch application process using procedures
  - Download patches directly from My Oracle Support and apply
  - Schedule to apply patches on multiple targets simultaneously
- **Configuration update:** After applying the patch, the configuration is updated with the latest and reports can be generated based on the applied patches.

**See Also:** "Managing Critical Patch Advisories" in the Grid Control online help and *Oracle Enterprise Manager Advanced Configuration*.

### Applying Critical Patches in Offline Mode

The Critical Patch Facility enables administrators to simply download the Critical Patch metadata from My Oracle Support and upload it to the repository. This metadata can then be used by the "RefreshFromMyOracleSupport" job for performing

Critical Patch calculations in offline mode. Administrators will be alerted to security updates—even if the Management Service is not connected to My Oracle Support.

Some data centers are not connected to the outside world. The Critical Patch Facility's offline mode makes it easy to keep your environment patched to the latest level. Subsequent patching can be done in offline mode as well, using the Software Library infrastructure.

### Accessing Critical Patch Advisories Pages in Grid Control

To access the Critical Patch Advisories pages in Grid Control:

- Click the **Deployments** tab, then click the link for the number of Patch Advisories in the Critical Patch Advisories for Oracle Homes section.
- Navigate to Grid Control Home page, then click the link for the number of Patch Advisories in the Critical Patch Advisories for Oracle Homes section.

This takes you to the Patch Advisories page, where you can view advisories, patch sets, and patches to apply, as well as affected Oracle homes and available "remedies."

## Application Server Patching

The patching application automates the deployment of Oracle patches for the application server and Management Agents. The application takes care of appropriate shutdown and startup of services and also allows execution of pre and post patching scripts to serve different use cases. Such flexibility makes mass deployment of interim patches and patchsets feasible even in complex multi-tier environments.

## Linux Host Patching

The "Patch Linux Hosts" tool, a powerful new feature in Grid Control, facilitates the automated management of Linux hosts in an enterprise. Use this feature to keep the Linux hosts in your enterprise up to date with vital software updates from your Linux vendor.

Patch Linux Hosts uses a reference-based grouped patching model, where you can create one or more reference package repositories containing up-to-date versions of various packages, and associate a group of Linux hosts with these package repositories.

The Patch Linux Hosts tool uses package repositories to patch the hosts as well as to monitor the deviation of the packages installed on the hosts. You can create different groups suited to your administrative needs and even associate different package repositories with different priorities for each group. You can independently control when and how often to update the hosts in the group, and how to determine their compliance with respect to the package repositories.

---

---

**Note:** To use this feature, make sure you have the following:

- Licenses for the Provisioning and Patch Automation Pack
  - Linux Management Pack
  - "Operator" privileges on the host that you want to patch
  - Ability to do sudo to the root user
- 
- 

The Linux patching feature provides the following functionalities:

- Setting up and managing RPM Repositories by subscribing to Unbreakable Linux network (ULN) channels
- Setting up and managing custom RPM Repositories and channels (cloning channels, copying packages from one channel into another, and deleting channels)
- Setting up Linux Patching Group to update a group of Linux hosts and compliance reporting from the Linux Patching group
- Scheduling Patching for non-compliant groups
- Managing Configuration file channels (creating/deleting channels, uploading files, and copying files from one channel into another)
- Patching through deployment procedures and emergency patching
- Undo Patching feature

### Enhanced Linux Patching for ULN

Enhanced Linux Patching feature of Enterprise Manager supports the Unbreakable Linux Network (ULN) subscribers through EM. ULN provides access to Linux software patches, updates and fixes for its customers. Oracle provides three levels of Unbreakable Linux support:

- Network Support - access to patches and updates via ULN
- Basic Support - access to patches and updates via ULN, 24x7 support, complete Linux server lifecycle management
- Premier Support - access to patches and updates via ULN, 24x7 support, Linux server lifecycle management, backporting, lifetime support

The Linux RPM Repository Server Setup page in Enterprise Manager allows you to set up a RPM repository server for Linux patching. You can select the Host to setup the RPM repository server and register the host to the Unbreakable Linux Network (ULN).

### Patching Features

- **Linux Host Patching Groups:** You can group a set of Linux hosts together to update all at once. Each group is associated with one or more package repositories that contain all the certified and appropriate versions of the software packages for the hosts of that group. Each group is configured with an update schedule for a recurring job to run to update the hosts with the associated package repositories.

**See Also:** "Creating a New Linux Host Group" in the Grid Control online help

- **RPM Repository:** RPM repository is a directory that contains RPM packages. The RPM repository is accessible via http or ftp. A RPM repository can be organized to contain packages from multiple channels.
- **Custom Channel:** A custom channel is a channel created by the user to store a set of custom RPM packages. Custom channels can be added to the RPM repository.
- **Configuration Channel:** A channel that is created by the user to store a set of Linux configuration files. Configuration channels can be used in the Linux patching application user interface to update configuration files on Linux hosts.
- **Compliance and automatic updates:** The compliance page contains information on the number of hosts in a group that are in compliance, as well as the number of "rogue" packages on a particular host. You can see metrics and charts to measure

compliance for all Linux Host Patching Groups, as well as historical compliance data.

- **Emergency Patching:** This feature gives you the option of performing "forced" updates, outside of the established schedule, to immediately respond to critical bugs or security alerts for all configured Linux hosts.
- **Undo Patching:** This feature adds flexibility by allowing you to roll back the software to its previous stable version, or even de-install the unstable version completely if that software version was found to be unsuitable or to have a bug or security vulnerability.
- **Patching through Deployment Procedures:** You can use deployment procedures to set up RPM repository, patch linux hosts, and perform other custom patching procedures.

## Operating Systems Patching

Enterprise Manager supports patching of Solaris, Linux, and Windows Operating Systems. For Solaris, you can directly connect to the vendor's Website and download patches.

In addition to proactive patching method mentioned in the section above, Enterprise Manager also supports ad-hoc patching of Linux, Windows and Solaris operating systems using native patching methods.

## Accessing Patching Pages in Grid Control

To access the patching pages in Grid Control:

- Click the **Deployments** tab, then click the links found under the Patching section:
  - Patching Through Deployment Procedures - This link takes you to the Deployment Procedure Manager page. Deployment procedures are best practices provided by Oracle for various Provisioning and Patching tasks. Procedures created by Oracle cannot be edited, but can be extended using 'Create Like', so that you can customize the procedure to fit your environment.
  - View/Upload Patch - This link takes you to the Patch Cache page and it provides a list of patches available in the Management Repository. These are patches that have been either automatically downloaded from My Oracle Support or manually uploaded to the patch cache.
  - Patch Linux Hosts - This link takes you to the Patch Linux Hosts page and helps keep Linux hosts up-to-date with vendor updates.
  - Patch Agent - This link takes you to the Agent Patching wizard.
- Click **Setup**, then click **Patching Setup** from the navigation pane. From this page, you can configure your settings for My Oracle Support and patching, proxy connection, offline patching, and Linux Staging server.





---

---

# Virtualization Management

This chapter introduces the concepts in virtualization management in the following sections:

- [Introduction to Virtualization Management](#)
- [Monitoring Oracle VM](#)
- [Administering Oracle VM](#)
- [Provisioning Oracle VM Server and Guest VM](#)

## Introduction to Virtualization Management

Virtualization management involves the monitoring, administration, and maintenance of virtual servers and guest virtual machines in your enterprise. Oracle provides the customer increased value by extending Enterprise Manager capabilities to monitor virtualized entities alongside the physical infrastructure and perform complete lifecycle management of virtual servers and software running on them.

---

---

**Note:** Refer to Support Note 781879.1 to enable Oracle VM Management functionality in Enterprise Manager Console.

Oracle recommends that no more than 24 Oracle VM servers be registered and managed with one Oracle Management Agent.

Enterprise Manager supports only Oracle VM server version 2.1.2 or above. All Oracle VM servers version 2.1.2 or above managed by Enterprise Manager should be updated to have OVS agent of version 2.2-70 or higher.

---

---

With Enterprise Manager, you receive:

- A complete set of integrated features for managing Oracle VM
- Unparalleled scalability that lets you manage a single Oracle VM or hundreds of instances
- An intuitive management product that is easy to use and deploy

---

---

**Note:** Oracle VM servers are remotely monitored by Enterprise Manager. One agent can monitor multiple virtual servers. The Enterprise Manager agent monitoring Oracle VM servers must be of version 10.2.0.5.0. Enterprise Manager agent must not be installed on virtual servers.

---

---

Following are commonly-used terms in virtualization:

### **Oracle VM**

Oracle VM is a server virtualization software. Installing Oracle VM on a bare metal box allows for the creation of multiple guest virtual machines on it. Just like Oracle databases and Oracle Fusion Middleware, Oracle VM is available for download on Oracle Technology Network (OTN).

### **Virtual Server**

Virtual Server is a generic term used to describe a physical box running virtualization software (hypervisor) on it. A new virtual server can be provisioned by installing Oracle VM server software on a bare metal physical box. "Virtual Server" is a target type in Enterprise Manager that represents the Oracle VM targets. Every Oracle VM server can perform one or more of the functions described below:

Every Oracle VM server can perform one or more of the functions described below:

- **Master Server function**

The Master Server is the core of the server pool operations. It acts as the contact point of the server pool to the outside world, and also as the dispatcher to other servers within the server pool. The load balancing is implemented by the Master Server. For example, when you start a guest virtual machine, the Master Server will choose a Guest VM Server with the maximum resources available to run the guest virtual machine. There can be only one Master Server in a server pool.

- **Utility Server function**

The Utility Server is responsible for I/O intensive operations such as copying or moving files. Its function focuses on the creation and removal operations of guest virtual machines, virtual servers, and server pools. There can be one or more Utility Servers in a server pool. When there are several Utility Servers, Server Pool Master chooses the Utility Server with the maximum CPU resources available to conduct the task.

- **Guest VM Server function**

The primary function of the Guest VM Server is to run guest virtual machines, thus acting as a hypervisor. There can be one or more Guest VM Servers in a server pool. When there are several Guest VM Servers, Master Server chooses the Guest VM Server with the maximum resources available (including memory and CPU) to start and run the virtual machine.

When the Oracle VM Server Pool is created, the user is asked to provide the details of the Master Server for that pool. By default, this Oracle VM Server also performs the functions of the Utility Server and Guest VM Server.

The user can later change the Utility Server and Guest VM Server functions using the Edit Virtual Server action.

### Monitoring Server

The monitoring server monitors the virtual server remotely. Multiple virtual servers are monitored by one agent. Enterprise Manager agent must not be installed on virtual servers.

### Virtual Server Pool

A Server Pool is a logical grouping of one or more virtual servers that share common storage. A virtual server can belong to only one virtual server pool at a time. Guest virtual machines and resources are also associated with server pools. Oracle VM Server Pool is an aggregate target type in Enterprise Manager to represent the server pool of Oracle VM Servers.

### Guest Virtual Machine

Guest Virtual Machine (also known as Guest VM) is the container running on top of a virtual server. Multiple guest virtual machines can run on a single virtual server. A guest virtual machine can be PXE-booted to install an operating system on it.

Similar to a physical machine with an operating system installed on it, a guest virtual machine with the operating system installed on it is termed as a 'Host' in Enterprise Manager. For example, a guest virtual machine running Oracle Enterprise Linux is a Linux Host.

A guest virtual machine monitored by an Enterprise Manager agent is treated like any other Host target. This type of guest virtual machine is referred to as a "Managed Guest VM".

A guest virtual machine that is not monitored by any Enterprise Manager agent is referred to as an "Unmanaged Guest VM". Only limited monitoring and configuration information is available for such a guest virtual machine.

## Virtualization Central Page as the Management Hub

The Enterprise Manager Virtualization Central page ([Figure 14–1](#)) shows the following:

- A quick view of the status of the virtual server and guest virtual machines and basic information about them
- CPU and Memory utilization of virtual servers
- Information about the server pool to which the virtual servers are registered
- Drilldowns to server pools, virtual servers, and guest virtual machines home pages from where you can view more information about these targets
- Allows you to perform administrative tasks such as starting and stopping guest virtual machines
- Allows you to perform provisioning activities such as creating virtual server pools, registering virtual servers, and creating guest virtual machines

Starting from the Virtualization Central page, you can access additional details and management functions by drilling down through the user interface. In addition, the Virtualization Central page provides a list of Related Links. From these links you can perform activities such as bare metal provisioning of Oracle VM server software, customization of table columns, and deployment of Enterprise Manager agent on a guest virtual machine.

**Note:** The Monitoring Server agent (Enterprise Manager agent that monitors Oracle VM servers) should be upgraded to version 10.2.0.5.0.

Figure 14–1 Virtualization Central Page

The screenshot displays the Oracle Enterprise Manager 10g Virtualization Central page. The page title is "Virtual Servers" and it includes a navigation bar with tabs for Home, Targets, Deployments, Alerts, Compliance, Jobs, and Reports. Below the navigation bar, there are search and action controls. The main content is a table with the following columns: Select, Name, Type, Status, Server Pool, Virtual Server, Operating System, Alerts, Policy Violation, Compliance Score (%), CPU Util (%), Memory Util (%), CPUs, and Memory Size (MB). The table contains five rows of data:

Select	Name	Type	Status	Server Pool	Virtual Server	Operating System	Alerts	Policy Violation	Compliance Score (%)	CPU Util (%)	Memory Util (%)	CPUs	Memory Size (MB)
<input type="checkbox"/>	agentVM1	Guest VM	Running	ca-emhost3-server-pool1	ca-emhost3.us.oracle.com	Enterprise Linux AS release 4 (October Update 5)	0 1	31 2 0	76	12.05	73.22	1	256
<input type="checkbox"/>	agentVM1	Guest VM	Running	ca-emhost3-server-pool1	ca-emhost3.us.oracle.com	Enterprise Linux AS release 4 (October Update 5)	0 0	4 0 0	82	12.05	47.25	1	256
<input type="checkbox"/>	ca-emhost3-server-pool1	Virtual Server Pool	Up										
<input type="checkbox"/>	ca-emhost3.us.oracle.com	Virtual Server	Up	ca-emhost3-server-pool1		Oracle VM server release 2.1.2	0 0	0 0 0	Unavailable	.7	55.24	2	2004
<input type="checkbox"/>	guestvm1	Guest VM	Editing	ca-emhost3-server-pool1	ca-emhost3.us.oracle.com					.2		3	256

Below the table, there are "Related Links" for Bare-Metal Provisioning, Customize Table Columns, and Deploy Agent. The footer contains copyright information and a link to "About Oracle Enterprise Manager".

A guest virtual machine can have the following additional status:

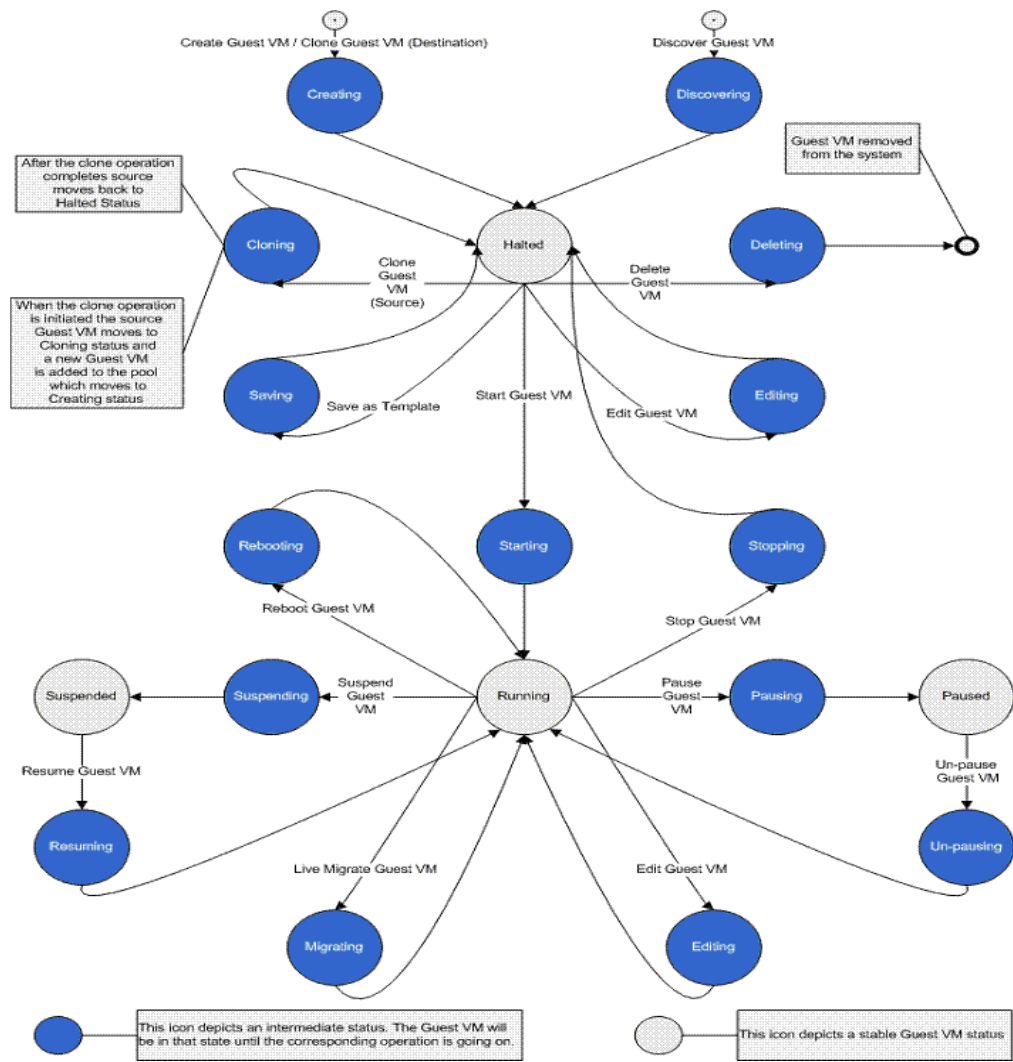
1. Operation Error - This is an intermediate status. The guest virtual machine moves to this status when an operation fails. The user has to perform the Clear Operation Error action to move the guest virtual machine out of this status.
2. Error - This status indicates that there are some errors with the guest virtual machine, and it cannot be used.
3. Unknown - This status indicates that Enterprise Manager has not been able to collect the state of the guest virtual machine.

**Table 14–1 Guest Virtual Machine Actions**

<b>Action</b>	<b>Applicable Status</b>	<b>Multi-Select/Single-Select</b>
Live Migrate Guest VM	Running	Multi-select
Clone Guest VM	Halted	Single-select
Save Guest VM as Template	Halted	Single-select
Edit Guest VM	Halted/Running	Single-select
Update Preferred Server List	Halted	Single-select
Destroy Guest VM	Halted/Error/Unknown	Multi-select
Start Guest VM	Halted/Error/Unknown	Multi-select
Stop Guest VM	Running/Paused/Error/Unknown	Multi-select
Reboot Guest VM	Running	Multi-select
Pause Guest VM	Running	Multi-select
Un-pause Guest VM	Paused	Multi-select
Suspend Guest VM	Running	Multi-select
Resume Guest VM	Suspended	Multi-select
VNC Console	Running	Single-select
Clear Operation Error	Operation Error	Single-select

[Figure 14–2](#) shows the guest virtual machine state transition.

**Figure 14–2 Guest Virtual Machine State Transition Diagram**



## Oracle VM Server Home Page

On the Oracle VM server home page (Figure 14–3), you can view basic information about the Oracle VM server including status, availability, operating system details, alerts, and guest virtual machines running on the Oracle VM server.

Following are the graphs described on the Oracle VM server home page:

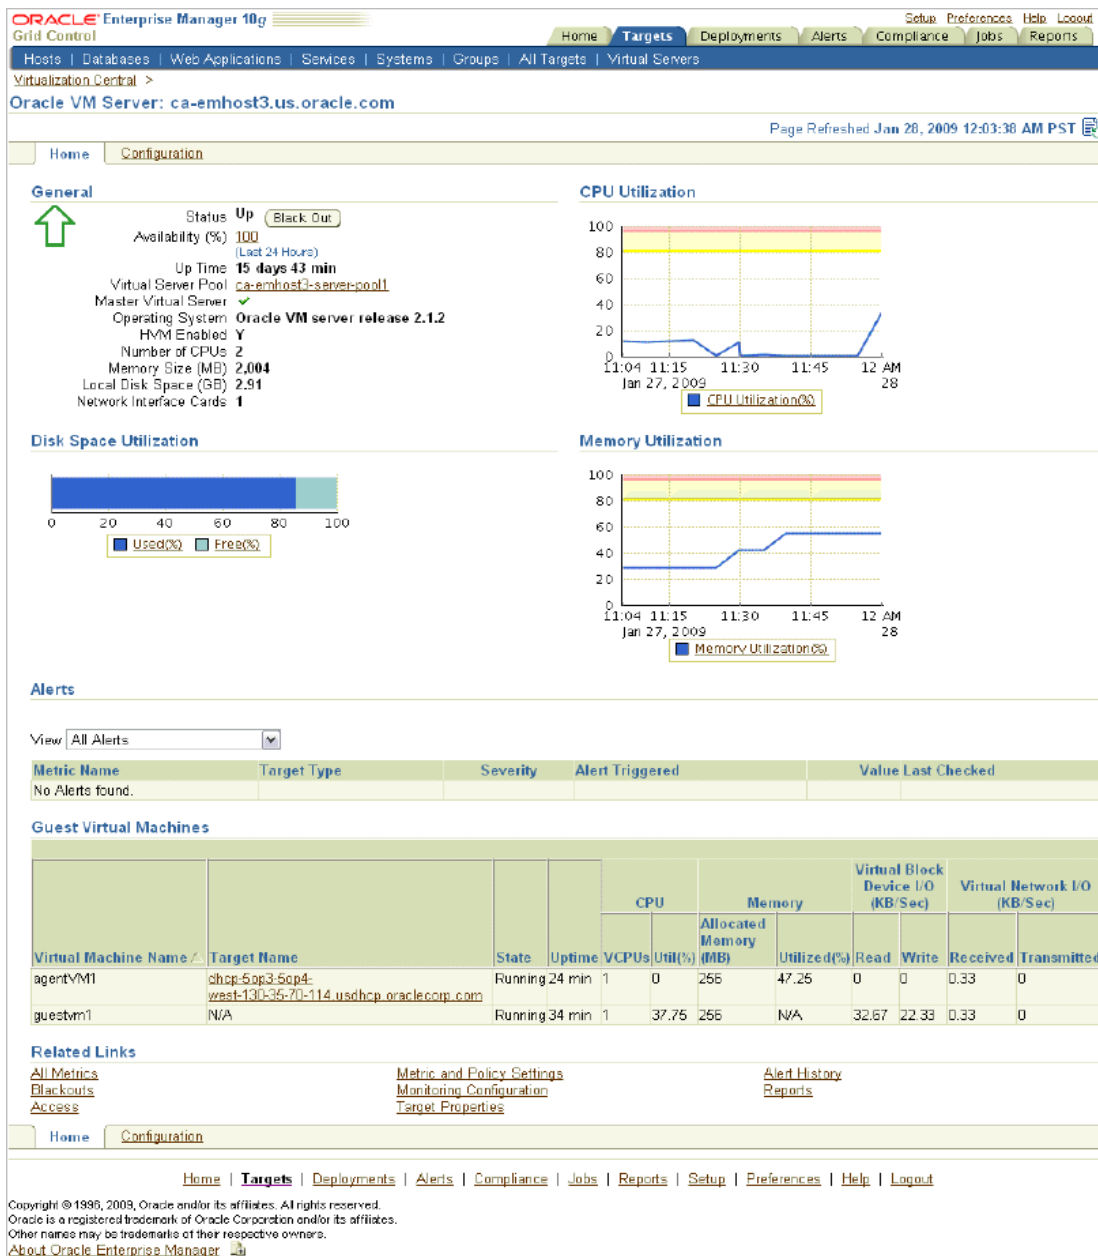
- The CPU Utilization graph provides information about the total physical CPU utilization(%) of the Oracle VM server.
- The Memory Utilization graph provides information about the physical memory utilization(%) of the Oracle VM server.
- The Disk Space Utilization graph provides information about the free and used local disk (partitioned) space on the Oracle VM server.

The Guest Virtual Machine section shows the following metrics for each of the guest virtual machines running on this Oracle VM Server.

- Virtual Machine Name - The name of the guest virtual machine running on the Oracle VM Server.

- Target Name - The Host target name if this is a managed guest virtual machine.
- State - State of the guest virtual machine. The state can be either "Running" or "Paused".
- Uptime - Displays how long this guest virtual machine has been "Running". This also includes the time this machine has been in the "Paused" state.
- CPU | vCPUs - Displays the number of Virtual CPUs this guest virtual machine has.
- CPU | Util(%) - Displays the amount of CPU utilization with respect to Total Physical CPU Utilization.
- Memory | Allocated - Displays the amount of memory allocated to this guest virtual machine.
- Memory | Util(%) - Displays the amount of memory utilized with respect to the memory allocated to the guest virtual machine.
- Virtual Block Device I/O (KB/sec) | Read - Displays the amount of data read from the block device.
- Virtual Block Device I/O (KB/sec) | Write - Displays the amount of data written to the block device.
- Virtual Network I/O (KB/sec) | Received - Displays the amount of data received by the Virtual Network interfaces.
- Virtual Network I/O (KB/sec) | Transmitted - Displays the amount of data transmitted by the Virtual Network interfaces.

Figure 14–3 Virtual Server Home Page



From the home page, you can drill down to more detailed performance information.

## Guest VM Home Page

The Guest Virtual Machine Home page is available only for managed guest virtual machines. It is similar to the Host Home page in Enterprise Manager. On this page, you can view basic information about the guest virtual machine including configuration details, alerts, policy violations, job details, patch advisories details, and security details. For more information about hosts, see [Host and Third-Party Target Management](#).

To get the complete benefits of managing and monitoring guest virtual machines using Enterprise Manager, the guest virtual machine must have an agent installed on it.



Figure 14–4 Guest Virtual Machine Home Page

The screenshot shows the Oracle Enterprise Manager 10g Grid Control interface. The main navigation bar includes Home, Targets, Deployments, Alerts, Compliance, Jobs, and Reports. The sub-navigation bar shows Hosts, Databases, Web Applications, Services, Systems, Groups, All Targets, and Virtual Servers. The host name is stapn10.us.oracle.com, and the latest data was collected on Jan 28, 2009 at 12:07:28 AM PST.

The page is divided into several sections:

- General:** Shows the status as Up, with a 'Black Out' button. The last boot was before 2 days 11 hours on Jan 25, 2009 at 1:00:12 PM (UTC-08:00). There are 2 logons, 100% availability (last 24 hours), and the guest virtual machine is 18\_agentVM1 on the ca-ermhost3.us.oracle.com server.
- Configuration:** Lists system details: Operating System (Enterprise Linux Enterprise Linux AS release 4), Hardware Platform (686), IP Address (152.68.100.181), CPUs (1), Memory Size (256 MB), and Local File Systems (3.81 GB).
- Alerts:** A table shows one alert: 'Filesystem Space Available (%) for /' with a severity of warning. The message indicates that the filesystem has 15.79% available space, falling below a warning (20%) or critical (5%) threshold. The alert was triggered on Jan 26, 2009 at 10:58:09 PM and last checked on Jan 27, 2009 at 7:35:23 PM.
- Policy Violations:** Shows 312 distinct policies violated, 110 compliance score, and 76% policy trend overview.
- Security:** Shows the last security evaluation on Jan 27, 2009 at 1:20:50 AM PST with a 68% compliance score.
- Critical Patch Advisories for Oracle Homes:** Shows 0 current advisories. A note states that patch advisory information may be stale because Oracle Support credentials are not configured.
- Job Activity:** Shows 1 problem execution, 0 action required executions, and 0 suspended executions. There are also 0 scheduled and 0 running executions.
- Related Links:** A list of links including Access, Alert History, All Metrics, Blackouts, Deployments, Execute Host Command, Log File Alerts, Metric and Policy Settings, Monitoring Configuration, Net Services Administration, Open Telnet Session, Privilege Delegation Setting, Remote File Editor, Reports, Storage Details, Target Properties, and User-Defined Metrics.

## Monitoring Oracle VM

Enterprise Manager provides comprehensive monitoring to enable you to identify the problem areas in your virtualization environment.

Enterprise Manager automatically gathers and evaluates diagnostic information from Oracle VM targets distributed across the enterprise. As with all targets managed by Enterprise Manager, an extensive number of Oracle VM metrics are automatically monitored against predefined thresholds. Alerts are generated in Grid Control when metrics exceed these thresholds.

## Administering Oracle VM

Oracle Enterprise Manager effectively keeps your Oracle VM available and running efficiently. Specifically, it provides a graphical user interface for managing Oracle VM Server, guest virtual machines, and server pools. As with Oracle VM monitoring,

administration of Oracle VM begins with the Virtualization Central page. You can perform administration tasks such as the following:

- Create, remove, and manage virtual server pools
- Register, deregister, edit, reboot, and stop virtual server
- Start, stop, suspend, resume, pause, un-pause, and edit guest virtual machines

## Editing Guest VM

A guest virtual machine can be edited as shown in [Figure 14–5](#).

**Figure 14–5 Edit Guest Virtual Machine Page**

ORACLE Enterprise Manager 10g  
Grid Control

Home Targets Deployments Alerts Compliance Jobs Reports  
Setup Preferences Help Logout

Hosts | Databases | Middleware | Web Applications | Services | Systems | Groups | All Targets | Virtual Servers  
Virtualization Central >

Edit Guest Virtual Machine: 60\_vncUr1f

Cancel Continue

General Disk Storage Network Interface Advanced

Review and modify the configuration of the Guest Virtual Machine.

\* Virtual Machine Name

Status **Running**

Server Pool Name **oexen4Pool**

Hosting Virtual Server **oexen-server-4.idc.oracle.com**

\* Memory (MB)

\* Dynamic Memory (MB)   
Dynamic Memory size must be atleast 256 MB

\* Virtual CPUs   
For a running Virtual Machine, increasing the number of cpus from the value at the last start, will take effect on next restart.

Description

\* VNC Console Password

\* Confirm VNC Console Password

General Disk Storage Network Interface Advanced

Cancel Continue

You can edit the Name, Memory, Dynamic Memory, vCPUs, VNC Console Password, add and remove disk storage, add network interfaces, and modify the boot order for the guest virtual machine. [Table 14–2](#) lists the guest virtual machine fields that can be edited and the action needed for the changes to take effect.

**Table 14–2 Guest Virtual Machines: Editable Fields**

Editable Field	Action for Changes to Take Effect
Virtual Machine Name	This field can be edited only for guest virtual machines in the Halted status. Changes will take effect immediately when the guest virtual machine is started.
Memory	This field can be edited only for guest virtual machines in the Halted status. Changes will take effect immediately when the guest virtual machine is started.

**Table 14–2 (Cont.) Guest Virtual Machines: Editable Fields**

Editable Field	Action for Changes to Take Effect
Dynamic Memory	For para-virtualized guest virtual machines, changes to the memory take affect without rebooting the guest virtual machine. For hardware-virtualized guest virtual machines, changes to memory take affect only after rebooting the guest virtual machine.
Virtual CPUs	For para-virtualized guest virtual machines, number of vCPUs can be edited and changes take affect without rebooting the guest virtual machine. For hardware-virtualized guest virtual machines, changes to vCPUs take affect only after rebooting the guest virtual machine.
VNC Console Password	New password takes affect after the guest virtual machine is stopped and restarted.
Add/Remove Disk Storage	Changes take place immediately.
Add/Remove Network Interface	Changes take place immediately.
Boot Order	Requires the guest virtual machine to be stopped and restarted.

If you want to restrict certain guest virtual machines to run on only certain Oracle VM servers, use the Preferred server list feature. Preferred server list contains all the Oracle VM servers on which this guest virtual machine is allowed to run. If the "Automatic" option is selected when creating a guest virtual machine, Enterprise Manager picks the Oracle VM server with maximum resources to run the guest virtual machine. When auto-restarting guest virtual machines during fail-over scenario, Enterprise Manager picks a preferred Oracle VM server to auto-restart the guest virtual machine.

## Provisioning Oracle VM Server and Guest VM

Enterprise Manager allows you to provision virtual servers and guest virtual machines in your virtualization enterprise.

You can do the following provisioning activities from the Virtualization Central page:

- Creating guest virtual machines (from Oracle VM template, ISO, PXE)
- Cloning guest virtual machines
- Live migrating guest virtual machines
- Saving guest virtual machine as a template
- Discovering guest virtual machines, templates, and ISO images
- Bare metal provisioning of Oracle VM server

You can also use deployment procedures to perform these provisioning operations.

### Creating Virtual Server Pool

For Enterprise Manager to manage Oracle VM servers, you will need to create a virtual server pool. Each virtual server pool must have at least one Oracle VM virtual server. The first virtual server added to the pool acts as the master server for the server pool.

[Figure 14–6](#) shows the Create Virtual Server Pool page. You can specify a virtual server pool and the master server.

**Figure 14–6 Create Virtual Server Pool Page**

**ORACLE Enterprise Manager 10g**  
Grid Control

Home Targets Deployments Alerts

Hosts | Databases | Middleware | Web Applications | Services | Systems | Groups | All Targets | Virtual Servers

Virtualization Central >

### Create Virtual Server Pool

Provide details of the virtual server pool and the master virtual server you want to register with Enterprise Manager.

\* Virtual Server Pool Name   
 Shared Directory   
 High Availability Mode  Enable

#### Master Virtual Server Details

\* Virtual Server Host   
Name of master Virtual Server

\* Monitoring Server Agent   
Grid Control Agent that will monitor the Virtual Server. Eg. abc.xyz.com:4889

\* Monitoring Server User   
Username of Grid Control Agent's Owner

\* Monitoring Server Password   
Password of Grid Control Agent's Owner

\* SSH Username   
SSH Username to connect to Virtual Server

\* SSH Password   
SSH Password to connect to Virtual Server

\* Oracle VM Agent Password   
OVM Agent Password on Virtual Server

\* OVS Proxy Location   
Staging location on Virtual Server

You will need to specify the following fields:

- Virtual Server Host — As part of the virtual server pool creation, one virtual server needs to be specified, which will act as the Master Server for the server pool.
- Monitoring Server Agent — Virtual servers in a server pool are monitored remotely by an Enterprise Manager agent. Specify the agent that will monitor the virtual servers.
- Monitoring Server User — Operating system user that owns the Enterprise Manager agent home/agent install and has read and write permission on agent home directory.
- Monitoring Server Password — Password for the user that owns the Enterprise Manager agent home/agent install.
- SSH Username and Password — User name and password for Enterprise Manager to SSH into the virtual server machine.
- Oracle VM Agent Password — Every Oracle VM server has an in-built Oracle VM agent running on it. Specify the password for this agent.
- OVS Proxy Location — This is a directory on the virtual server where the scripts required for monitoring are stored when the virtual server is registered with Enterprise Manager.

## Registering Virtual Servers

After creating a virtual server pool, you can register the virtual servers with the server pool (Figure 14–7). To register virtual servers, you will need to specify the fields described in [Creating Virtual Server Pool](#).

Figure 14–7 Register Virtual Server Page

ORACLE Enterprise Manager 10g  
Grid Control

Home Targets Deployments Alerts Compliance Jobs Reports

Hosts | Databases | Middleware | Web Applications | Services | Systems | Groups | All Targets | Virtual Servers

Virtualization Central >

### Register Virtual Server

Add details of the Virtual Servers to be registered. Cancel OK

**Batch Assignment** Assign

Enter values in input fields and click on Assign button to update the same in below table

Virtual Server Host: <input type="text" value="Virtual_server1"/> <small>Name of the Virtual Server.</small>	Monitoring Server Agent: <input type="text" value="dhcp-10-143-183-157.us.oracle.com"/> <small>GC Agent that will monitor the Virtual Server</small>
Monitoring Server Username: <input type="text" value="testt_user"/> <small>Username of GC Agent's owner</small>	Monitoring Server Password: <input type="password" value="*****"/> <small>Password of GC Agent's owner</small>
SSH Username: <input type="text" value="example_user"/> <small>SSH Username to connect to Virtual Server</small>	SSH Password: <input type="password" value="*****"/> <small>SSH Password to connect to Virtual Server</small>
Oracle VM Agent Password: <input type="password"/> <small>Password of OVM agent running on the Virtual Server</small>	OVS Proxy Location: <input type="text" value="staging_loc"/> <small>Staging location on Virtual Server</small>
Virtual Server Pool: <input type="text" value="apool"/> <small>Server Pool that this Virtual Server will join</small>	

Remove Duplicate | Add Another Row Add From File |

Select All | Select None

Select	Virtual Server Host	Monitoring Server Agent	Monitoring Server Username	Monitoring Server Password	SSH Username	SSH Password	Oracle VM Agent Password	OVS Proxy Location
<input type="checkbox"/>	Virtual_server1	dhcp-10-143-183-157.us.oracle.com	testt_user	*****	example_user	*****		staging_loc

## Creating Guest Virtual Machines

You can create guest virtual machines from the following three sources:

- Using Oracle VM Templates
- Using ISO images of Linux and Windows operating systems
- PXE bootable guest virtual machines that can boot over a network

Oracle VM template or ISO binaries can be available in any of the following locations:

- In the Software Library as a software component
- In the virtual server pool as a resource
- In a HTTP or NFS accessible external location

**Figure 14–8 Create Guest Virtual Machines Wizard**

ORACLE Enterprise Manager 10g  
Grid Control

Home Targets Deployments Alerts Compliance Jobs Reports  
Hosts | Databases | Middleware | Web Applications | Services | Systems | Groups | All Targets | Virtual Servers

Server Pool Source Configure Schedule / Credentials Review

Create Guest Virtual Machines: Server Pool

Cancel Step 1 of 5 Next

\* Number of Virtual Machines   
Total number of virtual machines to be created

\* Server Pool

Preferred Server List  Automatic  Manual

Figure 14–8 shows the Create Guest Virtual Machines: Server Pool page from the guest virtual machine creation wizard. Navigate to the Create Guest Virtual Machines: Configure page and specify the following fields:

- Virtual Machine Name Prefix — This string that will be used as a prefix in the guest virtual machine name. When more than one guest virtual machine is created using the Enterprise Manager Console, guest virtual machine names will be prefix1, prefix2, and so on.
- Root Password — This is the operating system "root" password to log into the guest virtual machine once it is up and running.
- Number Of Cores — Number of CPU cores to assign to this guest virtual machine. This will show up as virtual CPU (vCPU) configuration parameter for the guest virtual machine.
- Memory Size (MB) — The amount of memory to be allocated to the guest virtual machine. The physical virtual server machine should have sufficient memory to run the guest virtual machine. Otherwise, the guest virtual machine will be created, but will not be started.
- VNC Console Password — Enterprise Manager provides a VNC console to the newly created guest virtual machine. This is the VNC password to access the VNC console.
- Number of Network Interfaces — Number of virtual network interfaces this guest virtual machine should have. The minimum value is 1.
- Start VM after creation — Specify whether the guest virtual machine should be started automatically after it is created.
- Enable High Availability Mode — When creating the guest virtual machine, if the "Auto-restart" flag is selected, this guest virtual machine will participate in the High Availability functionality. If its parent virtual server fails, this guest virtual machine will be auto-restarted on another Oracle VM server.

Enabling High Availability mode for the virtual server pool allows all the Oracle VM servers in the pool to be clustered. This provides for high availability/fail-over support. For more information about high availability failover for server pools, see *Oracle VM Server User's Guide* available at:

[http://download.oracle.com/docs/cd/E11081\\_01/doc/doc.21/e10898/ha.htm#CHDEFABI](http://download.oracle.com/docs/cd/E11081_01/doc/doc.21/e10898/ha.htm#CHDEFABI)

To enable High Availability, you will need the following configuration:

- High Availability mode should be set for the virtual server pool.
- Clustering configuration should be done for Master Server and other virtual servers in the server pool.
- Auto-restart flag must be set for the guest virtual machines.
- Additional Disks — Every guest virtual machine is created with at least one disk. Specify additional disks and their sizes. These disks will be added to the guest virtual machine. After the guest virtual machine is created, log into the guest virtual machine and reformat and mount the additional disks.
- Shared Disks — Oracle VM allows sharing of disks between multiple guest virtual machines. Specify a shared disk that the guest virtual machine can access.





---

---

## Oracle Beehive Service Dashboard Metrics

This appendix details the Oracle Beehive Service Dashboard metrics.

### A.1 Oracle Beehive Metrics

[Table A-1](#) explains the Oracle Beehive Service Dashboard metrics collected by Enterprise Manager:

**Table A-1 Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
Client Service	Performance	Connect Time	Web Transaction	Connect Time (ms)	Oracle Beehive Zimbra client is a web-based client to access E-mail, Calendar, etc., Test for Zimbra client has been modeled as a web transaction. Enterprise Manager breaks down each transaction into individual phases. Performance metrics for each phase of the transaction can help you pinpoint the cause of a slow response time alert. Connect Time is the total time spent in the transaction connecting to the server. There may be multiple connections made during a transaction. Time spent waiting for requests that result in redirects count as Redirect Time rather than Connect Time.

**Table A-1 (Cont.) Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
Client Service	Performance	Perceived Time	Web Transaction	Perceived Time per Page (ms)	The average amount of time that it would take the beacon to play each page in the transaction. This is a good metric for setting thresholds because it is the closest active measurement of what the user-experience is likely to be. Because it is normalized on a per-page basis, Perceived Time per Page is also a good metric for comparing the relative performance of different transactions.
Conferencing Service	Usage	Conf Participants	Conference Application	Conference Participants	A Conference Session can have multiple participants. Conf Participants is the participants taking part in conference sessions at each collection.

**Table A-1 (Cont.) Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
Conferencing Service	Usage	Conf Sessions	Conference Application	Conference Sessions	Oracle Beehive provides online voice and Web conferencing enabling meeting organizers and participants to conduct collaborative sessions online through Oracle Beehive Conferencing. There can be multiple conference sessions in progress at the same time. Conf Sessions is the number of conference sessions active at each collection (current conference sessions).

**Table A-1 (Cont.) Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
Email Service	Performance	Email Read Time	IMAP Test	Time to Read Email (ms)	Oracle Beehive E-mail Service supports standard SMTP protocol to send an e-mail and IMAP protocol to read e-mails. Msg Delivery Time measures the round-trip time (in milliseconds) from when a sender sends an e-mail to when the same e-mail is received in the receiver's Inbox. It indicates the time taken to deliver an email. This is a SMTP service test metric and is collected only if the Calculate Message Delivery Time test parameter in SMTP test is set to Yes.
Email Service	Performance	IMAP Login Time	IMAP Test	Login Time (ms)	Oracle Beehive E-mail Service supports standard IMAP protocol to access e-mail. IMAP Login Time is the time taken (in milliseconds) to login to an IMAP Service.

**Table A-1 (Cont.) Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
Email Service	Performance	Msg Delivery Time	SMTP Test	Message Delivery Time (ms)	Oracle Beehive E-mail Service supports standard SMTP protocol to send an e-mail and IMAP protocol to read e-mails. Message Delivery Time measures the round-trip time (in milliseconds) from when a sender sends an e-mail to when the same e-mail is received in the receiver's Inbox. It indicates the time taken to deliver an email. This is a SMTP service test metric and is collected only if the Calculate Message Delivery Time test parameter in SMTP test is set to Yes.
Email Service	Usage	Emails Received	Email Application	Transport Mails Received	The number of e-mails received by the Oracle Beehive E-mail Service transport subsystem in a collection interval.
Email Service	Usage	IMAP Sessions	Email Application	Current IMAP Sessions	Number of E-mail Service IMAP sessions at each collection.

**Table A-1 (Cont.) Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
Email Service	Usage	SMTP Sessions	Email Application	Current Open SMTP Sessions	Number of E-mail service SMTP sessions at each collection.
FTP Service	Performance	Download Time	FTP Test	Download Time (ms)	The total time required to download a file using Oracle Beehive FTP service.
FTP Service	Performance	Login Time	FTP Test	Login Time (ms)	Oracle Beehive FTP Service supports all content management-related features, including download and upload, through FTP and FTP over TLS. Login Time is the time required to login to the Oracle Beehive FTP service.
FTP Service	Performance	Upload Time	FTP Test	Upload Time (ms)	The time required to upload a file to Oracle Beehive server using Oracle Beehive FTP service.
FTP Service	Usage	Files Downloaded	FTP Application	Downloaded Files	Number of files downloaded using Oracle Beehive FTP Service in a collection interval.
FTP Service	Usage	Files Uploaded	FTP Application	Uploaded Files	Number of files uploaded using Oracle Beehive FTP Service in a collection interval.

**Table A-1 (Cont.) Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
FTP Service	Usage	Users	FTP Application	User FTP Connections	Number of FTP connections in a collection interval.
Mobile Service	Performance	Login Time	PushMail Application	Avg Login Time	The Oracle Beehive Mobile Mail Service provides a complete Push IMAP (P-IMAP) v0.6 implementation for real-time delivery of e-mail to users' mobile devices. Avg Login Time is the average login time to login to Oracle Beehive PushMail service as measured in each collection interval.
Mobile Service	Usage	Emails Read	PushMail Application	Emails Marked Seen	Number of e-mail messages marked as seen in a collection interval that were accessed using Mobile Mail (PushMail) Service.
Mobile Service	Usage	Emails Sent	PushMail Application	Emails Sent	Number of e-mail messages sent in a collection interval using Mobile Mail (PushMail) Service.
Mobile Service	Usage	Users	PushMail Application	Current Server Usage	Number of users using Mobile Mail Service (PushMail) at each collection.



**Table A-1 (Cont.) Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
Presence and IM Service	Usage	Contact Listing Time	XMPP Test	Contacts Listing Time (ms)	Time taken (in milliseconds) to retrieve the roster of contacts from Oracle Beehive XMPP Service.
Presence and IM Service	Usage	Files Transferred	XMPP Application	File Transfers	Number of files transferred in a collection interval using Oracle Beehive XMPP Service.
Presence and IM Service	Usage	IM Sending Time	XMPP Test	Message Send Time (ms)	Time taken (in milliseconds) to send an instant message using Oracle Beehive XMPP Service.
Presence and IM Service	Usage	Messages Exchanged	XMPP Application	IM Messages Exchanged	Number of messages exchanged in a collection interval using Oracle Beehive XMPP Service.
Presence and IM Service	Usage	XMPP Login Time	XMPP Test	Login Time (ms)	Oracle Beehive XMPP Service supports Instant Messaging using standard XMPP protocol. XMPP Login Time is the time taken (in milliseconds) to login to the Oracle Beehive XMPP Service.
Presence and IM Service	Usage	XMPP Sessions	XMPP Application	Active XMPP Sessions	Number of active XMPP sessions running as measured in a collection interval.

**Table A-1 (Cont.) Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
Time Management Service	Performance	Add Meeting Time	CalDAV Test	Event Add Time (ms)	Time taken (in milliseconds) to add a calendar event using CalDAV Service.
Time Management Service	Performance	CalDAV Login Time	CalDAV Test	Calendar Access Verification Time (ms)	Oracle Beehive CalDAV and Time Management Services support all aspects of time and task management. CalDAV Login Time is the time taken (in milliseconds) to access and login to Oracle Beehive Calendar using CalDAV URL.
Time Management Service	Performance	Content Listing Time	CalDAV Test	Content Listing Time (ms)	Time taken (in milliseconds) to list the contents of the Calendar folder using Beehive CalDAV Service.
Time Management Service	Usage	Meetings Created	TimeManagement Application	Number of Create Meetings	Number of meetings created in Oracle Beehive Calendar in a collection interval.
Time Management Service	Usage	Meetings Searched	TimeManagement Application	Number of Search Meetings	Number of meetings searched in Oracle Beehive Calendar in a collection interval.

**Table A-1 (Cont.) Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
Time Management Service	Usage	Tasks Created	TimeManagement Application	Number of Create Tasks	Number of tasks created in Oracle Beehive Calendar in a collection interval.
User Directory Service	Performance	Create User Time	UserDirectory Application	Avg Create User Time	Oracle Beehive User Directory Service (UDS) provides comprehensive user and group management capabilities for Oracle Beehive. Create User Time is the average time taken to create a user in Oracle Beehive user directory as measured in a collection interval.
User Directory Service	Performance	Update AddrBook Time	UserDirectory Application	Avg Update Address Book Time	Average time taken to update an address book in Oracle Beehive user directory as measured in a collection interval.
User Directory Service	Performance	Update User Time	UserDirectory Application	Avg Update User Time	Average time taken to update a user in Beehive user directory as measured in a collection interval.
User Directory Service	Usage	AddrBooks Updated	UserDirectory Application	Number of Address Books Updated	Number of Oracle Beehive User Directory address books updated in a collection interval.

**Table A-1 (Cont.) Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
User Directory Service	Usage	Users Created	UserDirectory Application	Number of Users Created	Number of users created in Oracle Beehive User Directory in a collection interval.
User Directory Service	Usage	Users Updated	UserDirectory Application	Number of Users Updated	Number of users updated in Oracle Beehive User Directory in a collection interval.
Voice Message Service	Performance	Create VM Time	VoiceMessage Application	Avg Create Voice Message Time	Oracle Beehive Voice Message Service supports all aspects of voicemail and fax management for Oracle Beehive. Create VM Time is the average time taken to create a Voice Message using Oracle Beehive Voice Message Service as measured in a collection interval.
Voice Message Service	Performance	Delete VM Time	VoiceMessage Application	Avg Delete Voice Message Time	Average time taken to delete a Voice Message using Oracle Beehive Voice Message Service as measured in a collection interval.
Voice Message Service	Usage	VMs Created	VoiceMessage Application	Number of Create Voice Messages	Number of Oracle Beehive voice messages created in a collection interval.

**Table A-1 (Cont.) Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
Voice Message Service	Usage	VMs Deleted	VoiceMessage Application	Number of Delete Voice Messages	Number of voice messages deleted in a collection interval.
WebDAV Service	Performance	Connect Time	WebDAV Test	Connect Time (ms)	Oracle Beehive WebDAV Service supports all content management-related features, including download and upload, using the Web-based Distributed Authoring and Versioning (WebDAV) protocol. Connect Time is the time taken (in milliseconds) to connect to the Beehive WebDAV server.
WebDAV Service	Performance	File Creation Time	WebDAV Test	File Creation Time (ms)	Indicates time taken (in milliseconds) to create the file using Oracle Beehive WebDAV Service.
WebDAV Service	Performance	File Read Time	WebDAV Test	File Read Time (ms)	Indicates time taken (in milliseconds) to read the file from Beehive using WebDAV protocol.

**Table A-1 (Cont.) Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
WebDAV Service	Usage	Files Downloaded	WebDAV Application	Number of GET Commands	Number of GET commands executed using Oracle Beehive WebDAV Service in a collection interval.
WebDAV Service	Usage	Files Uploaded	WebDAV Application	Number of PUT Commands	Number of PUT commands executed using Oracle Beehive WebDAV Service in a collection interval.
WebDAV Service	Usage	Users	WebDAV Application	Number of Logins	Number of Logins operations performed to Oracle Beehive WebDAV Service in a collection interval.
Workspaces Service	Performance	Create Doc Time	Workspaces Application	Avg Create Document Time	Oracle Beehive Workspace Service supports all the features and functionality provided by Oracle Beehive personal and team workspaces. Create Doc Time is the average time taken to create a document in Oracle Beehive Workspace as measured in a collection interval.

**Table A-1 (Cont.) Oracle Beehive Metrics**

<b>Service Name</b>	<b>Metric Type</b>	<b>Metric Label</b>	<b>Source - Test/Target Type</b>	<b>Source Metric Name</b>	<b>Description</b>
Workspaces Service	Performance	Create Team Ws Time	Workspaces Application	Avg Create Team Workspace Time	Average time taken to create a team workspace in Oracle Beehive as measured in a collection interval.
Workspaces Service	Performance	Delete Doc Time	Workspaces Application	Avg Delete Document Time	Average time taken to delete a document in Oracle Beehive Workspace as measured in a collection interval.
Workspaces Service	Usage	Docs Created	Workspaces Application	Number of Create Documents	Number of documents created in Oracle Beehive Workspaces in a collection interval.
Workspaces Service	Usage	Docs Deleted	Workspaces Application	Number of Delete Documents	Number of documents deleted in Oracle Beehive Workspaces in a collection interval.
Workspaces Service	Usage	Team Ws Created	Workspaces Application	Number of Create Team Workspaces	Number of team workspaces created in Oracle Beehive in a collection interval.





---

---

# Glossary

## **adaptive alert thresholds**

Computed statistical alert thresholds using either static (user-defined) or dynamic (self-adjusting) baselines.

*See also* [metric baselines](#).

## **ADDM**

*See* [Automatic Database Diagnostic Monitor \(ADDM\)](#).

## **administrator account**

Enterprise Manager administrator account that provides permission to perform daily administration work and access administrative information.

## **agent**

*See* [Oracle Management Agent](#).

## **alert**

Indicator signifying a particular metric condition has been encountered. An alert is triggered when one of the following conditions is true:

- A metric threshold is reached.
- The availability of a monitored service changes. For example, the availability of the host changes from up to down.
- A metric-specific condition occurs. For example, an alert is triggered whenever an error message is written to a database alert log file.

## **Application Server Control**

Web-based management features designed to monitor and administer application server instances, farms, and clusters.

## **attestation**

Process of having users or system managers confirm people's access rights on a periodic basis.

## **Automatic Database Diagnostic Monitor (ADDM)**

Application that automatically takes regularly scheduled snapshots of the database activity. ADDM identifies the most resource-intensive components or operations and provides advice, which may recommend running an advisor or making configuration changes to your database.

### **Automatic Storage Management (ASM)**

Application that automates and simplifies the layout of data files, control files, and log files. Database files are automatically distributed across all available disks, and database storage is rebalanced whenever the storage configuration changes.

### **Automatic target discovery**

Process by which targets are located and added to Enterprise Manager. Automatic discovery begins when the Oracle Management Agent starts up after installation. Targets located on the Agent are discovered and added to Enterprise Manager to be monitored and administered. How the targets are installed determines which targets are automatically added.

### **Automatic Workload Repository (AWR)**

Automatic capture of statistics data for real-time and historical performance analysis. AWR includes snapshot data, active session history data, and workload reports.

### **availability**

The percentage or amount of scheduled time that a computing system provides application service.

### **AWR**

See [Automatic Workload Repository \(AWR\)](#).

### **Beacon**

Application to monitor transactions from different user communities or geographical regions. You can enable the Beacon software to monitor the availability and performance of network components (a host computer or an IP traffic controller) from more than one location, or Web pages and Web applications from multiple network locations.

### **blackout**

Maintenance tool for suspending any data collection activity on one or more monitored targets. This allows you to perform scheduled maintenance on targets while excluding these special-case situations from the data analysis to obtain a more accurate, long-term picture of a target's performance.

### **cluster cache coherency**

For Oracle Real Application Clusters (RAC) environments, helps you identify processing trends and optimize performance. The Cluster Cache Coherency page enables you to view the global cache block access latency, global cache block transfer rate, and global cache block transfers and physical reads.

### **control**

See [Oracle Enterprise Manager 10g Grid Control](#), [Oracle Enterprise Manager 10g Database Control](#), [Oracle Enterprise Manager 10g Application Server Control](#), [Oracle Enterprise Manager 10g Application Server Control for Collaboration Suite](#).

### **dashboard**

Tool for proactively monitoring status and alerts as they occur by viewing the health of managed targets in real time. The dashboard presents information using intuitive icons and graphics that let you spot recent changes and quickly identify and respond to problems.

**database configuration**

The configuration information Enterprise Manager collects for an Oracle database, which includes:

- General database and instance information, such as the database name, instance name, and whether or not the database is running in restricted mode or archive log mode
- Initialization parameter values
- System Global Area values
- Tablespaces and their parameters
- Datafiles and their parameters
- Control files and their attributes
- Redo logs and their attributes
- Rollback segments and their parameters
- High availability information
- License information
- Database options information

**database instance**

Running Oracle Database consisting of memory structures (SGA) and background processes. An instance only exists while it is up and running. Essentially, a database resides on disk, while an instance resides in memory.

A database is normally managed by only one instance. However, when using Oracle Real Application Clusters, multiple instances can be started for a single database (on different hosts of a cluster).

*See also* [Oracle Real Application Clusters \(RAC\)](#).

**enterprise configuration**

Configuration information stored in the Management Repository for the set of hosts and targets that comprise your enterprise.

**fast-start failover**

Ability of Data Guard to rapidly and automatically fail over to a standby database without requiring manual intervention. This improves the degree of availability as well as the disaster resilience for the system infrastructure.

**fetchlet**

Parameterized data access mechanism that takes arguments for input and returns formatted data.

**Grid Control console**

*See* [Oracle Enterprise Manager 10g Grid Control console](#).

**group**

Single logical unit that can include targets of the same type (for example, all your production databases) or include targets of different types (for example, all targets comprising your business's application). Groups enable you to enable you to collectively monitor and administer many targets.

### **home page**

See [target home page](#).

### **host configuration**

Configuration information that Enterprise Manager collects for a host that is a managed target, including:

- Hardware for the host (including memory, CPU, I/O device, and network interface information)
- Operating system for the host (including information such as operating system properties, packages, and installed patches)
- Installed Oracle software, including (but not limited to) installed products and their components, patch sets, and interim patches on the host. Enterprise Manager uses the Oracle Universal Installer inventory or inventories on a host to obtain information about the Oracle products installed on the host.

### **instance**

See [database instance](#).

### **J2EE**

Abbreviation for Java 2 Platform Enterprise Edition. J2EE is an environment for developing and deploying enterprise applications.

### **Management Agent**

See [Oracle Management Agent](#).

### **Management Plug-in**

See [Oracle Management Plug-in](#).

### **Management Repository**

See [Oracle Management Repository](#).

### **Management Services**

See [Oracle Management Service](#).

### **metric**

Unit of measurement used to report the health of the system.

### **metric baselines**

Named period of time associated with a target and used as a reference for evaluating target performance. Statistics are computed over the baseline period for specific target metrics. You can use these statistics to automatically set metric thresholds for alerting, as well as to normalize graphical displays of system performance.

See also [adaptive alert thresholds](#).

### **metric snapshot**

Named collection of a target's performance metrics that have been collected at a specific point in time. A metric snapshot can be used as an aid in calculating metric threshold values based on the target's past performance.

**Notification Methods**

Mechanism for defining different mechanisms for sending notifications. These include e-mail, SNMP traps, or running custom scripts—or all three. Once defined, these methods can then be used with Notification Rules for sending notifications to administrators as a result of alert occurrences.

**Oracle ecosystem**

Organization of software, hardware, and their environment so that they function as a unit. Typically, this consists of an Oracle platform and all third-party software, including storage systems, hosts, routers, and so on. The Grid Control Home page provides a unified view of your ecosystem.

**Oracle Enterprise Manager 10g**

Oracle Enterprise Manager is the Oracle integrated management solution for managing your computing environment.

**Oracle Enterprise Manager 10g Application Server Control**

Application Server Control relies on various underlying technologies to discover, monitor, and administer the Oracle Application Server environment. Application Server Control consists of the Application Server Control console and its underlying technologies:

- Oracle Dynamic Monitoring Service (DMS)
- Oracle Process Management Notification (OPMN)
- Distributed Configuration Management (DCM)
- A local version of the Oracle Management Agent specifically designed to gather monitoring data for Application Server Control.

**Oracle Enterprise Manager 10g Application Server Control console**

The Enterprise Manager Web-based user interface for managing Oracle Application Server 10g. The Application Server Control console is installed and available with every Application Server 10g installation.

From the Application Server Control console, you can monitor and administer a single Oracle Application Server instance, a farm of Oracle Application Server instances, or Oracle Application Server Clusters.

**Oracle Enterprise Manager 10g Application Server Control for Collaboration Suite**

Application Server Control for Collaboration Suite consists of a Web-based user interface and a set of underlying technologies. Oracle Enterprise Manager 10g Application Server Control for Collaboration Suite is installed and available with every Collaboration Suite installation.

**Oracle Enterprise Manager 10g Application Server Control Console for Collaboration Suite**

The Enterprise Manager Web-based user interface for managing Oracle Collaboration Suite. The Application Server Control for Collaboration Suite console is installed and available with every Oracle Collaboration Suite installation. From this console, you can monitor and administer a single Oracle Collaboration Suite instance.

### **Oracle Enterprise Manager 10g Database Control**

Database Control relies on various underlying technologies to discover, monitor, and administer the Oracle Database environment.

Database Control consists of the Database Control console and its underlying technologies:

- A local version of the Oracle Management Service designed to work with the local database or clustered database
- A local Oracle Management Repository installed in the local database and designed to store management data for the Database Control console
- A local version of the Oracle Management Agent designed to provide monitoring data to the local Management Service

### **Oracle Enterprise Manager 10g Database Control console**

The Enterprise Manager Web-based user interface for managing Oracle Database 10g. The Database Control console is installed and available with every Oracle Database 10g installation.

From the Database Control console, you can monitor and administer a single Oracle Database instance or a clustered database.

### **Oracle Enterprise Manager 10g Grid Control**

Grid Control relies on various underlying technologies to discover, monitor, and administer your computing environment. Grid Control consists of the Grid Control console and these underlying technologies:

- One or more Oracle Management Services
- The Oracle Management Repository
- The central Management Agent
- Remote Oracle Management Agents, installed on each monitored host

### **Oracle Enterprise Manager 10g Grid Control console**

Enterprise Manager Web-based user interface for centrally managing your entire computing environment.

From the Grid Control console, you can monitor and administer your entire computing environment from one location on the network. All the services within your enterprise, including hosts, databases, listeners, application servers, HTTP Servers, and Web applications, are easily managed as one cohesive unit.

### **Oracle Management Agent**

A process deployed on each monitored host. The Oracle Management Agent is responsible for monitoring all targets on the host, for communicating that information to the middle-tier Management Service, and for managing and maintaining the host and its targets.

### **Oracle Management Plug-in**

Target type provided by the user or a third party to extend Enterprise Manager's set of predefined target types.

### **Oracle Management Repository**

Two tablespaces in an Oracle Database that contain information about administrators, targets, and applications that are managed within Enterprise Manager.

The Management Service uploads to the Management Repository the monitoring data it receives from the Management Agent. The Management Repository then organizes the data so it can be retrieved by the Management Service and displayed on Grid Control.

**Oracle Management Service**

A J2EE Web application that renders the user interface for the Grid Control console, works with all Management Agents in processing monitoring and job information, and uses the Management Repository as its data store.

**My Oracle Support**

Oracle Support Services site where customers can get information about released patches and outstanding bugs. You can use Enterprise Manager to download patches from *My Oracle Support*.

**Oracle Streams**

Propagation and management of data, transactions, and events in a data stream either within a database, or from one database to another. Controls what information is put into a stream, how the stream flows from database to database, what happens to messages in the stream, and how the stream terminates.

**policy**

Defines the desired behavior of systems and is associated with one or more targets or groups. Policies include different categories of policy rules, such as configuration, security, and storage rules. Enterprise Manager compares the targets for which policy rules exist with the policy rules for that target type, and identifies the policy violations for the target.

**policy association**

Relationship of a policy rule and its settings to a target or monitoring template.

**policy compliance score**

Number denoting the conformance of a target to a set of requirements (policy rules.) The score, ranging from 0% to 100%, is much like a test score. A score of 100% means that the target was fully compliant with the goals of the policy. The policy compliance score facilitates your assessment of the relative levels of attention needed for various targets given the violations on those targets.

**policy evaluation**

Process of testing a policy's condition and recording any violations in the repository.

**policy groups**

Standards that report deviations and enable closed loop remediation by optionally taking action to bring systems back into compliance.

**policy rule**

A conditional expression that tests values from a target against a condition; for example, verifying that database profile limits are set as expected.

**policy rule library**

Collection of policy rules.

**policy settings**

Conditions related to a policy. The policy settings include, but are not limited to, parameter values and excluded objects.

**policy trend**

The policy compliance trend of a policy in the domain of a particular target.

**policy violation**

Infringement of a policy rule. For example, one of the host security policy rules checks for open ports. The recommendation is that the insecure ports not be opened. Therefore, the violation is that there are open ports.

**privileges**

A right to perform management actions within Enterprise Manager, such as view any target and add any target in the enterprise, or a right to perform operations on a target such as maintaining and cloning the target. Types of privileges are defined by Oracle.

**Oracle Real Application Clusters (RAC)**

Option that allows multiple concurrent instances to share a single physical database. It provides a high-availability database environment spanning multiple hosts. Each cluster can be comprised of multiple cluster databases, each of which consists of multiple cluster database instances.

**redundancy group**

Group that contains members of the same type that function collectively as a unit.

**repository**

See [Oracle Management Repository](#).

**roles**

Collection of predefined Enterprise Manager target or system privileges created by super administrators to facilitate granting multiple privileges or roles to users. Roles limit target access and access to specific management features.

**Secure Socket Layer (SSL)**

Software to secure the communications between the Management Service and the Management Agent. Encrypted communications between the Management Agent and Management Service over HTTPS ensures the privacy of data sent from one computer to another.

**service**

Entity that provides a useful function to its users, such as CRM applications, online banking, and email services. Some simpler forms of services are business functions that are supported by protocols such as DNS, LDAP, POP, or SMTP.

**standby database**

Read-only database on the standby node. The standby database, also referred to as the replicated database, is physically identical to the master database. In the event of catastrophic failures, data modification activities “fail over” to the standby database so that it becomes the new master database.



**Super Administrator**

Enterprise Manager administrator that can create, modify, and delete any Enterprise Manager administrator, create any role in the system, perform any action on any target in the system, and see all areas of the management system.

**SYSMAN**

By default during the installation of Oracle Enterprise Manager, one Super Administrator account is created with the user name of SYSMAN. The SYSMAN account should be used to perform infrequent system-wide, global configuration tasks such as setting up the environment. Other administrator accounts can be created for daily administration work. The SYSMAN account is:

- Owner of the Management Repository schema
- Default Enterprise Manager Super Administrator
- User name used to log in to Enterprise Manager the first time

**system**

Logical grouping of targets that collectively hosts one or more services. It is a set of infrastructure targets (hosts, databases, application servers, and so on) that function together to host one or more applications or services.

**system privilege**

Allows a user to perform system-wide operations. For example, the VIEW ANY TARGET system privilege allows the administrator to view any target on the system, including Oracle Management Agents and Oracle Management Services.

**target**

A single component that you can monitor or configure with Enterprise Manager. Examples of a target include:

- Single Oracle 10g database
- Group of databases that provide your worldwide customers with product information
- Oracle Application Server or an instance of Oracle HTTP Server
- Web application that your customers visit to investigate or buy your products
- Linux host computer, including its memory, disks, and CPU
- Server load balancer switch that controls the Internet traffic for a set of Web servers

Enterprise Manager can manage all these targets. A complete list of the target types you can manage is included in *Oracle Enterprise Manager Grid Control Installation and Basic Configuration*.

**target compliance score**

Average of the individual policy compliance scores for policies that have been associated with a target. Also, a target's security compliance score is the average of the compliance scores for security policies associated with that target.

You can use a compliance score to assess a target's overall compliance, and a group can use a score to determine the most troubled targets with respect to policy violations.

**target home page**

Page that contains general information about the selected target. From a target home page, you can drill down for more detailed performance information.

**target privilege**

Privilege that allows an administrator to perform operations on a target. For example, the View Target privilege allows an administrator to view properties, inventory, and monitor information about a target.

**threshold**

Boundary values against which monitored metric values are compared. The comparison determines whether an alert should be generated.

**topology**

Graphically shows the relationships between components, nodes, or objects within different Oracle applications.

**wait classes**

Grouping of wait events. Whenever an Oracle process waits for something, it records the wait using one of a set of predefined wait events. These wait events are grouped in wait classes, which show how much of the database is waiting for CPU time. The wait classes appear in the Database Performance page.

## A

---

accessing configuration pages, 4-7  
Active Session History (ASH), databases, 5-5  
Active Sessions Waiting page, picture of, 5-6  
ADDM, 3-1  
    purpose of, 3-1  
    using, 5-8  
    viewing findings, 5-8  
administering  
    databases, 5-14  
    hosts, 9-9  
    Linux hosts, 9-9  
administrators  
    defining rules, 2-3  
    responsibilities today, 2-3  
Advisor Central Home page, 5-19  
advisors  
    database administration, 5-18  
    invoking, 5-18  
    versus alerts, 5-19  
Aggregate Service service type, 6-2  
alerts  
    Collaboration Suite, 8-30, 10-5, 11-8  
    initiating with AWR, 5-3  
    monitoring Oracle Application Server, 8-11  
    Oracle Application Server monitoring, 8-11  
    server-generated, 3-1  
    services, 6-9  
    versus advisors, 5-19  
Analyze feature, services and, 6-11  
application database administrator, role of in data  
    masking, 5-20  
Application Server  
    analyzing historical performance, 8-22  
    backing up, 8-22  
    cloning, 8-23  
    diagnosing performance issues, 8-21  
    extensible monitoring, 8-21  
    management, 8-1  
    managing configurations, 8-20  
    out-of-box management, 8-1  
    patching, 8-23  
    recovering, 8-22  
    topology, viewing, 8-12  
    viewing topology, 8-12

Application Server Control, 8-1  
    for Collaboration Suite, 10-7  
applying critical patches, 13-10  
automated  
    RAC deployments, picture of, 2-4  
Automatic Database Diagnostic (ADDM) page,  
    picture of, 5-9  
Automatic Database Diagnostic Monitor  
    *See* ADDM  
automatic restarts, RAC, 5-40  
Automatic Storage Management (ASM), 5-14  
Automatic Workload Repository (AWR)  
    databases and, 5-3  
availability  
    definition, 6-2  
    diagnosing Collaboration Suite problems, 8-30,  
        10-5, 11-8  
    in services, 6-2  
    of a service, 6-2  
    services dashboard, 6-7  
Average Active Sessions chart, databases, 5-4, 5-5

## B

---

backing up  
    Application Server environment, 8-22  
    databases, 5-32  
    logical database objects, 5-35  
backup settings, 5-34  
baselines  
    change management and, 5-18  
    metric, 5-12  
    services and, 6-9  
BEA WebLogic, 9-11  
    extensions to Grid Control, 2-6  
    picture of Server Cluster page, 2-6  
Beacons  
    Collaboration Suite, 8-28, 10-2, 11-5  
    services and, 6-2  
benefits of  
    Oracle Grid, 2-1  
    Oracle Grid Control, 2-1  
BIG-IP  
    *see* load balancer switches  
blocking sessions, hang analysis feature, 5-13  
business metrics, services and, 6-3

business needs and Grid Control, 2-1

## C

---

change management, 5-18

Client Configuration Collection Tag, 4-7

client configurations, 4-6

Client System Analyzer (CSA), 10-5

client configurations, 4-6

deployed independently, 4-7

in Grid Control, 4-6

cloning

Application Server, 8-23

databases, 5-35

operation, 13-4

production database in data masking, 5-21

Cluster Cache Coherency page, picture of, 5-38

cluster cache coherency, RAC, 5-37

cluster interconnects, 5-38

Cluster Interconnects page, picture of, 5-39

Collaboration Suite

administration for instances, 10-7

Application Server Control for, 10-7

Beacons and, 10-2

collecting and using configuration data, 10-5

configuring services, 10-2

deployment, 10-1

diagnosing problems, 8-30, 10-5, 11-8

Grid Control features for use with, 8-31, 10-7, 11-8

monitoring and alerts, 8-30, 10-5, 11-8

monitoring services, 8-29, 10-3, 11-7

reports, 10-6

Root Cause Analysis, 8-30, 10-5, 11-8

service topology, 10-2

services, 10-1

services dashboard, 10-2

systems, 10-4

Collaboration Suite Service Home page, picture of, 10-3

Collaboration Suite Service Topology page, picture of, 10-4

collected configurations for targets, 4-1

comparing configurations, 4-4, 5-36

configuration pages, accessing in Grid Control, 4-7

configurations

client, 4-6

Client System Analyzer (CSA), 4-6

comparing, 4-4, 5-36

features, databases, 5-14

hardware and software, 4-1

managing, 8-20

searching an enterprise, 4-5

using database configuration features, 5-14

viewing for an enterprise, 4-4

configuring

collected configurations for targets, 4-1

hardware and software, 4-1

services for Collaboration Suite, 10-2

software, reducing time with provisioning

tools, 2-3

controlling distribution of resources, 5-15

converting

dictionary-managed tablespace, 5-15

single instances to RAC, 5-15

crash recovery, 5-33

Create Service page, picture of, 5-40

creating

database forms, 6-11

masking definition, 5-20

systems, 6-4, 6-5

creating guest virtual machines, 14-13

critical patch advisories, 13-9

managing, 13-10

providing, 5-36

Critical Patch Facility, 13-7

Critical Patch metadata, downloading, 13-10

## D

---

dashboard

Collaboration Suite and, 10-2

definition and function, 2-2

picture of services dashboard, 2-2

services, 6-6

Data Guard, overview of, 5-34

data masking

Add Columns page, 5-27

Add Dependent Columns page, 5-30

adding columns to a masking definition, 5-23

application database administrator, role of, 5-20

Canadian social insurance numbers, 5-26

cloning the production database, 5-21, 5-30

Create Masking Definition page, 5-27

creating masking definition, 5-20

Data Masking Definitions page, 5-27

Define Column Mask page, 5-28, 5-31

Define Group Mask page, 5-29

defining new formats, 5-23

description, 5-19

DM\_FMTLIB package, installing, 5-26

exporting and importing definitions, 5-32

information security administrator, role of, 5-20

ISBN numbers, 5-26

major task steps, 5-21

mask format libraries, description of, 5-20

masking definitions, description of, 5-20

masking format templates, using, 5-24

North American phone numbers, 5-26

other Oracle security products, 5-32

patterns of format definitions, 5-25

post-processing functions, 5-24

predefined masking formats, using, 5-25

primary key, 5-20

recommended workflow, 5-21

Script Generation Results page, 5-29

security administrator, 5-20

social security numbers, 5-26

staging region, 5-21

UPC numbers, 5-26

- use cases, 5-31
- user-defined functions, 5-24
- workflow, 5-21
- data, hiding with data masking, 5-19
- Database Control and Grid Control, comparisons of, 5-1
- Database Home page
  - features and information shown, 5-2
  - picture of, 5-3
- Database Performance page, picture of, 5-4
- Database Resource Manager, 5-15
- databases
  - Active Session History (ASH), 5-5
  - ADDM, 5-8
  - administering, 5-14
  - administration, 5-14
  - advisors, 5-18
  - Automatic Storage Management (ASM), 5-14
  - Average Active Sessions chart, 5-4, 5-5
  - backing up, 5-32
  - backup and recovery settings, 5-34
  - change management, 5-18
  - cloning, 5-35
  - converting dictionary-managed tablespace, 5-15
  - converting single instance to RAC, 5-15
  - Data Guard overview, 5-34
  - Database Resource Manager, 5-15
  - diagnosing performance problems
    - automatically, 5-8
  - Export to Files feature, 5-35
  - flashback recovery, 5-33
  - hang analysis feature, 5-13
  - historical SQL, 5-10
  - home page, 5-2
  - Home Page capabilities, 5-2
  - Host chart, 5-5
  - Import from Database feature, 5-35
  - indexes and data retrieval, 5-16
  - instance activity, 5-10
  - Instance Disk I/O chart, 5-5
  - Instance Throughput chart, 5-5
  - Load Data from File feature, 5-35
  - maintaining, 5-32
  - management, introduction, 5-1
  - managing
    - materialized views, 5-17
    - objects, 5-16
    - restore points, 5-34
    - users and privileges, 5-17
  - Memory Access Mode, 5-13
  - metric baselines, 5-12
  - Monitor Export and Import Job feature, 5-35
  - monitoring, 5-3
  - optimizing performance, 5-14
  - Oracle Streams, 5-35
  - patching software, 5-36
  - recovery, 5-33
  - Recovery Manager (RMAN), 5-32
  - scheduling jobs, 5-16
  - schemas, working with, 5-16

- Segment Advisor, 5-16
- services, cluster-managed, 5-39
- session details, viewing, 5-7
- software patching, 5-36
- SQL details, viewing, 5-6
- top activity, viewing, 5-8
- top consumers of system resources, 5-9
- Topology Viewer, 5-11
- tracking statistics to improve performance, 5-16
- transporting tablespaces, 5-35
- using configuration features, 5-14
- views, 5-16
- wait class drill-downs, 5-5
- XML features, 5-17
- Dell PowerEdge Linux hosts, 9-7
- deploying software, reducing time with provisioning tools, 2-3
- deployments
  - applying critical patches, 13-10
  - Client Configuration Collection Tag, 4-7
  - Client System Analyzer (CSA), 4-7
  - Collaboration Suite, 10-1
  - Critical Patch Facility, 13-7
  - Linux host patching, 13-11
  - Management Repository, 4-1
  - managing critical patch advisories, 13-10
  - managing for hosts, 9-11
  - patching, 13-7
  - patching Oracle software, 13-8
- diagnosing
  - Collaboration Suite problems, 8-30, 10-5, 11-8
  - performance issues, Application Server, 8-21
  - performance problems automatically, 5-8
  - performance problems interactively, 5-5
  - service problems, 6-12
  - slow or hung systems, 5-13
  - Web application problems, 6-12
- diagnostics
  - Request Performance Diagnostics feature, 6-13
- dictionary baseline, 5-18
- dictionary-managed tablespace, converting, 5-15
- directory monitoring, 9-7
- Distributed Configuration Management (DCM), 8-2, 10-7
- documentation, getting from OTN, 3-2
- downloading Critical Patch metadata, 13-10

---

**E**

- EMC SMARTS integration with Grid Control, 2-6
- emergency patching, 13-13
- End-User Performance Monitoring feature, 6-10
- enterprise configuration, 4-4
  - hardware, 4-4
  - predefined searches, 4-5
  - searching, 4-5
  - user-defined searches, 4-6
  - viewing, 4-4
- Execute Host Command page, description of, 9-10
- Export to Files feature, 5-35

- exporting files, 5-35
- extended application server management, 2-1
- extending Enterprise Manager
  - lifecycle stages for plug-ins, 3-5
- extensible monitoring, 8-21

## F

---

- F5
  - see* load balancer switches
- figures
  - Active Sessions Waiting Page, 5-6
  - Automated RAC Deployments and Extensions, 2-4
  - Automatic Database Diagnostic Monitor Page, 5-9
  - Cluster Cache Coherency Page, 5-38
  - Cluster Interconnects Page, 5-39
  - Collaboration Suite Service Home Page, 10-3
  - Collaboration Suite Service Topology Page, 10-4
  - Consolidated Dashboard View of Multiple Services, 6-7
  - Create Service Page, 5-40
  - Database Home Page, 5-3
  - Database Performance Page, 5-4
  - Historical SQL Page, 5-11
  - Host Configuration Page, 4-3
  - Host Hardware Details Page, 9-6
  - Host Home Page, 9-3
  - Host Operating System Details Page, 9-5
  - Host Performance Page, 9-4
  - Host Storage Details Page, 9-9
  - Instance Activity Page, 5-10
  - Metric Baselines Page, 5-13
  - Monitoring BEA WebLogic, 2-6
  - Monitoring the Performance of Your Deployed J2EE Applications, 8-12
  - Monitoring Top JSPs for a J2EE Application, 8-22
  - Monitoring Top Servlets for a J2EE Application, 8-22
  - Oracle System Monitoring Dashboard, 8-14
  - Searching for Patches, 2-5
  - Service Topology Page, 6-9
  - Services Dashboard Showing Multiple Services, 2-2
  - Session Details Page, 5-7
  - SQL Details Page, 5-7
  - System Home Page, 6-6
  - System Topology Page, 6-8
  - Top Consumers Top Modules Page, 5-10
  - Topology for SQL Details Plan, 5-12
- file monitoring, 9-7
- flashback recovery, 5-33
- forms transactions, 6-11

## G

---

- Generic Service service type, 6-2
- Grid Control
  - ability to eliminate costs and errors, 2-3
  - accessing configuration pages, 4-7

- administrators and, 2-3
- application server management, 8-2
- Client System Analyzer, 4-6
- compared with Database Control, 5-1
- extensions and plug-ins, 2-6
- functionality in Release 2, 2-2
- leveraging for Collaboration Suite, 8-31, 10-7, 11-8
- log file monitoring, 9-6
- lower management costs and, 2-2
- managing host deployments, 9-11
- monitoring hosts, 9-2
- third-party management, 9-11
- Guaranteed Restore Point, 5-34
- guest virtual machine, 14-3
- guest VM server, 14-2

## H

---

- hang analysis, for hung systems, 5-13
- hardware
  - configuration, collecting information, 4-1
  - enterprise configuration, 4-4
  - monitoring, Dell PowerEdge, 9-7
- hiding data using data masking, 5-19
- High Availability Oracle Clusterware, 5-40
- Historical SQL page, picture of, 5-11
- historical SQL, databases, 5-10
- Host chart, databases, 5-5
- Host Configuration page
  - description of, 9-4
  - picture of, 4-3
- Host Hardware Details page
  - description of, 9-5
  - picture of, 9-6
- Host Home page
  - description of, 9-2
  - picture of, 9-3
- host management feature, 9-1
- Host Operating System Details page
  - picture of, 9-5
- Host Performance page
  - description of, 9-3
  - picture of, 9-4
- Host Storage Details page
  - description of, 9-8
  - picture of, 9-9
- hosts
  - administration, 9-9
  - Dell PowerEdge Linux, 9-7
  - Execute Host Command page, description of, 9-10
  - executing host commands remotely, 9-10
  - Host Configuration page
    - description of, 9-4
  - Host Hardware Details page
    - description of, 9-5
  - Host Home page, description of, 9-2
  - Host Performance page, 9-3
  - Host Storage Details page, description of, 9-8

- Linux host patching, 13-11
- Linux hosts administration, 9-9
- managing deployments, 9-11
- monitoring, 9-2
- Operating System Details page
  - description of, 9-4
- Program Resource Utilization page, 9-6
- Remote File Editor page, 9-10
- remote host target, 9-10
- running operating system commands for, 9-10
- storage resource tracking, 9-8
- HP OpenView Operations integration with Grid Control, 2-6
- HTTP Server session volumes, 8-11

## I

---

- IBM WebSphere, 9-11
  - extensions to Grid Control, 2-6
- Import from Database feature, 5-35
- importing files, 5-35
- indexes, data retrieval performance and, 5-16
- Information Publisher
  - patches and, 13-9
- information security administrator, role of in data masking, 5-20
- installing software, reducing time with provisioning tools, 2-3
- Instance Activity page, picture of, 5-10
- instance activity, databases, 5-10
- Instance Disk I/O chart, databases, 5-5
- Instance Throughput chart, databases, 5-5
- integration with non-Oracle solutions, 2-6
- Interactive Transaction Tracing, slow performance and, 6-12
- IP traffic controllers, 9-13
- IT infrastructure performance, 2-1

## J

---

- J2EE application responsiveness, monitoring, 8-11
- JBoss, 9-11
- Job System for hosts, 9-10
- jobs
  - Job System for hosts, 9-10
  - predefined tasks, 3-4
  - purpose of, 3-3
  - scheduling, 5-16

## L

---

- Linux
  - host patching, 13-11
  - host patching groups, 13-12
- load balancer switches
  - F5 Networks, 9-14
  - F5 server load balancers, extensions of, 2-6
- Load Data from File feature, 5-35
- log file monitoring, 9-6

## M

---

- maintaining databases, 5-32
  - backups, 5-32
  - Data Guard, 5-34
  - flashback recovery, 5-33
  - managing restore points, 5-34
  - recovery, 5-33
- management
  - Oracle Application Server instances, 8-2
  - out-of-box using application server control, 8-1
  - third-party, 9-11
- Management Plug-ins
  - extensions and, 2-6
  - introduction, 2-6
- Management Repository, 4-1
  - deployments and, 4-1
- managing
  - backups, 5-33
  - configurations, 8-20
  - critical patch advisories, 13-10
  - database objects, 5-16
  - databases, Home page, 5-2
  - hosts, 9-1
  - materialized views, 5-17
  - patches, 2-4
  - restore points, 5-34
  - storage objects, 5-14
  - systems, 6-4
  - third-party targets, 9-11
  - users and privileges, 5-17
- mask format libraries, 5-20
- masking definitions, 5-20
  - credit card numbers, 5-25
- master server, 14-2
- materialized views, managing, 5-17
- media recovery, 5-33
- Memory Access Mode, performance issues and, 5-13
- metric baselines, 5-12
- Metric Baselines page, picture of, 5-13
- metrics
  - metric baselines, 5-12
  - performance and usage of services, 6-3
- Micromuse integration with Grid Control, 2-6
- modeling business applications using services, 6-2
- Monitor Export and Import Job feature, 5-35
- monitoring
  - business applications using services, 6-2
  - Collaboration Suite, 8-30, 10-5, 11-8
  - databases, 5-3
  - end-user performance, 6-11
  - extensible, Application Server, 8-21
  - files and directories, 9-7
  - hardware for Dell PowerEdge systems, 9-7
  - hosts, 9-2
  - log files, 9-6
  - Oracle Application Server, 8-11
  - Program Resource Utilization page, 9-6
  - RAC, 5-37
  - services, 6-6
  - services for Collaboration Suite, 8-29, 10-3, 11-7

- systems, 6-5
- templates for services, 6-4
- Web application services, 6-10

Monitoring Templates feature  
for services, 6-4

My Oracle Support  
advisories on patches

## N

---

NetApp Filers

- extensions of, 2-6

Network Appliance Filer, 9-13

new functionality in Grid Control Release 2, 2-2

notifications services, 6-9

## O

---

object privileges, 5-17

objects

- dictionary baseline, 5-18
- managing, 5-16

offline patching, 13-9

online patching, 13-9

OPatch, 4-5

Operating System Details page, description of, 9-4

optimizing

- database performance, 5-14
- performance in RAC, 5-37

Oracle

- Application Server
  - historical performance, 8-22
  - instances, 8-2
  - monitoring, 8-11
  - Web Cache, 8-11
- Clusterware, 5-40
- Clusterware Registry, 5-41
- Dynamic Monitoring Service (DMS), 8-2, 10-7
- Grid Control, benefits of, 2-1
- home credentials, patching and, 13-9
- Management Agent, 8-11
- Management Agent, configurations and, 5-36
- patching software, 13-8
- Process Management Notification (OPMN), 8-2, 10-7
- Streams, overview of, 5-35
- Technology Network (OTN), 3-2

Oracle Forms, 6-11

Oracle Grid Control

- See Grid Control

Oracle MetaLink

- See My Oracle Support

Oracle VM, 14-2

OTN and documentation, 3-2

Outlook Express mail clients, 10-6

out-of-box

- management using application server control, 8-1

## P

---

patch management, 2-4

patches, searching for, 2-5

patching

- Application Server, 8-23
- applying critical patches, 13-10
- compliance, 13-12
- Critical Patch Facility, 13-7
- deployment, 13-7
- emergency, 13-13
- information publisher, 13-9
- Linux host patching groups, 13-12
- Linux hosts, 13-11
- offline, 13-9
- online, 13-9
- OPatch, 4-5
- Oracle home credentials, 13-9
- Oracle software, 13-8
- setup, 13-9
- shared agent patching, 13-9
- software, 5-36
- standalone database patching, 13-9
- undoing, 13-13
- validation, 13-9

performance

- ADDM, 5-8
- analyzing historical, 8-22
- charts for services, 6-9
- diagnosing Collaboration Suite problems, 8-30, 10-5, 11-8
- diagnosing problems
  - automatically, 5-8
  - interactively, 5-5
  - with top reports, 8-21
- for services, 6-2
- hang analysis feature, 5-13
- historical SQL, 5-10
- Host Performance page, description of, 9-3
- increasing with indexes, 5-16
- instance activity, 5-10
- Interactive Transaction Tracing, 6-12
- Memory Access Mode, 5-13
- metric baselines, 5-12
- metrics, services and, 6-3
- optimizing, 5-14
- optimizing in RAC, 5-37
- Oracle Application Server, historical, 8-22
- service management and, 6-3
- services dashboard, 6-7
- session details, viewing, 5-7
- SQL details, viewing, 5-6
- top activity, viewing, 5-8
- top consumers of system resources, 5-9
- tracking statistics, 5-16
- wait class drill-downs, 5-5

plug-ins for management, 2-6

post-processing functions, data masking and, 5-24

predefined

- job tasks, 3-4
- searches, enterprise configuration, 4-5

primary key, data masking and, 5-20

privileges



- managing for users, 5-17
- Show Grantees function, 5-17
- provisioning Oracle VM, 14-11

## R

---

### RAC

- adding cluster hosts for efficiency, 5-37
- automatic restart of key components, 5-40
- cluster cache coherency, 5-37
- cluster interconnects, 5-38
- cluster-managed database services, 5-39
- converting single instances to RAC, 5-15
- interconnects, 5-38
- monitoring, 5-37
- optimizing performance, 5-37
- Oracle Clusterware, 5-40
- Oracle Clusterware Registry, 5-41
- purpose of, 5-37
- tasks, 5-37
- viewing current state of interfaces, 5-38
- Voting Disk, 5-40

Real Application Clusters

- See RAC

recovery

- Application Server environment, 8-22
- flashback, 5-33
- settings, 5-34
- types of, 5-33

Recovery Manager (RMAN), for backups, 5-32

regulatory compliance using masked data, 5-19

Remote File Editor page, 9-10

remote host target, 9-10

reports

- for services, 6-9
- Oracle Collaboration Suite, 10-6

Request Performance Diagnostics feature, 6-13

resource tracking, storage, 9-8

restarts, RAC, 5-40

restore points, managing, 5-34

Right to Financial Privacy Act of 1978, 5-20

Root Cause Analysis

- Collaboration Suite, 8-30, 10-5, 11-8
- for services, 6-2, 6-12
- service failure and, 6-8
- Topology Viewer, 6-12

running operating system commands, 9-10

## S

---

Sarbanes-Oxley requirements, 5-19

scheduling jobs, 5-16

schemas

- objects, materialized views, 5-17
- working with, 5-16

searching

- enterprise configuration, 4-5
- for patches, picture of, 2-5

security

- compliance with masked data, 5-19
- data masking, 5-19

- list of Oracle products, 5-32
- mask format libraries, 5-20
- masking definitions, 5-20
- security administrator, data masking and, 5-20
- Segment Advisor, for managing databases, 5-16
- server-generated alerts, 3-1
- Service Level Management feature, 6-4
- service test, 6-2
- Service Topology page, picture of, 6-8
- service types
  - Aggregate Service, 6-2
  - Generic Service, 6-2
  - Web Application, 6-2
- service, definition of, 6-1
- service-level agreements, 6-1
- services, 6-2
  - alerts, 6-9
  - Analyze feature, 6-11
  - availability of, 6-2
  - baselines, 6-9
  - Beacons, 6-2
  - business metrics, 6-3
  - cluster-managed database services, 5-39
  - Collaboration Suite, 10-1
  - configuring for Collaboration Suite, 10-2
  - dashboard, 6-6
    - Collaboration Suite, 10-2
  - dashboard, picture of, 2-2
  - diagnosing problems, 6-12
  - diagnosing Web application problems, 6-12
  - End-User Performance Monitoring feature, 6-10
  - failure and Root Cause Analysis, 6-12
  - forms and, 6-11
  - forms transactions, 6-11
  - Interactive Transaction Tracing, 6-12
  - introduction, 6-1
  - key Beacons, 6-3
  - key components, 6-3
  - key tests, 6-3
  - monitoring, 6-6
  - monitoring for Collaboration Suite, 8-29, 10-3, 11-7
  - Monitoring Templates feature, 6-4
  - monitoring templates for, 6-4
  - monitoring Web applications, Web application services, monitoring, 6-10
  - notifications, 6-9
  - performance and usage, 6-3
  - performance definition, 6-2
  - performance metrics, 6-3
  - performance of in charts, 6-9
  - reports for, 6-9
  - Root Cause Analysis, 6-12
  - Root Cause Analysis tool, 6-2
  - service level agreements, 6-1
  - service level definition, 6-2
  - Service Level Management feature, 6-4
  - service test, 6-2
  - service test-based availability, 6-3
  - Service Topology page, 6-8

- setting service-level rules, 6-4
- system-based availability, 6-3
- topology for Collaboration Suite, 10-2
- Topology Viewer, 6-5
- transactions and, 6-10
- usage definition, 6-2
- usage metrics, 6-3
- Web Application Management feature, 6-10
- Services Dashboard, 6-6
- session details, viewing, 5-7
- setting service-level rules, 6-4
- shared agent patching, 13-9
- Shared Global Area (SGA), for slow or hung systems, 5-13
- slow performance, Interactive Transaction Tracing and, 6-12
- software configurations
  - collecting information, 4-1
  - working with, 5-36
- Software Library, using, 13-6
- software patching, 13-8
- SQL
  - details, viewing, 5-6
  - historical, 5-10
- SQL Details page, picture of, 5-7
- stage and apply, for patches, 13-9
- staging region, data masking and, 5-21
- standalone database patching, 13-9
- standby databases, Data Guard and, 5-34
- storage
  - devices, 9-13
  - objects, managing, 5-14
  - resource tracking, 9-8
- System Home page, picture of, 6-6
- System Topology page, picture of, 6-7
- system-based availability, 6-3
- systems
  - Collaboration Suite, 10-4
  - creating, 6-4, 6-5
  - definition of, 6-4
  - managing, 6-4
  - monitoring, 6-5
  - privileges, 5-17
  - System Topology page, 6-7
  - topologies, viewing and creating, 6-5
- SYSUSER operations, 5-17

## T

---

- tablespaces, transporting, 5-35
- third-party
  - application servers, support for, 9-11
  - plug-ins, support for, 9-12
- third-party management, 9-11
  - BIG-IP load balancer switches, 9-14
  - EMC SMARTS, 2-6
  - HP OpenView Operations, 2-6
  - Micromuse, 2-6
  - Network Appliance Filer, 9-13
- Top Activity page, picture of, 5-8

- top activity, viewing, 5-8
- Top Consumers Top Modules page, picture of, 5-10
- top consumers, databases, 5-9
- Top J2EE Applications, 8-11
- topology
  - application server, 8-12
  - database, 5-11
  - service, 6-8
  - system, 6-7
- Topology for SQL Details Plan, picture of, 5-12
- Topology Viewer
  - application server, 8-12
  - database, 5-11
  - definition of, 6-5
  - Root Cause Analysis information, 6-12
  - service, 6-8
  - system, 6-7
- tracking
  - statistics for databases, 5-16
  - storage resource allocation, 9-8
- transporting tablespaces, 5-35

## U

---

- undoing patching, 13-13
- usage metrics, services and, 6-3
- use cases, data masking, 5-31
- user-defined
  - functions, data masking and, 5-24
  - searches, enterprise configuration, 4-6
- utility server, 14-2

## V

---

- viewing
  - enterprise configuration, 4-4
  - session details, 5-7
  - SQL details, 5-6
  - top activity, 5-8
- views, database objects and, 5-16
- virtual server, 14-2
- virtual server pool, 14-3
- Virtualization Central, 14-3
- Voting Disk, 5-40

## W

---

- wait classes, databases, 5-5
- Web Application Management feature, services and, 6-10
- Web Application service type, 6-2
- Workload Repository, database performance and, 5-16

## X

---

- XML database features, 5-17